

# Security Now! #707 - 03-26-19

## Tesla, Pwned

### This week on Security Now!

This week on Security Now! We have the return of "Clippy", Microsoft's much-loathed dancing paperclip, operation "ShadowHammer" which reports say compromised ASUS (... but did it?), the ransomware attack on Norsk Hydro aluminum, the surprise renaming of Windows Defender, a severe bug revealed in the most popular PDF generating PHP library, an early look at Microsoft's forthcoming Chromium-based web browser, hope for preventing caller ID spoofing, a needed update for users of PuTTY, Mozilla's decision to conditionally rely upon Windows' root store, Microsoft to offer virtual Windows 7 and 10 desktops through Azure, details of the Windows 7 End of Life warning dialog... then a bit of Sci-Fi, SQL and SpinRite news, followed by our look at the results of the much anticipated Mid-March Vancouver Pwn2Own competition... One of the results of which our episode title gives away!

It looks like you're  
sending out a  
ballistic missile  
alert.



## Security News

### Welcome back Clippy!

In a blog posting made on April 11, 2001 titled "Farewell Clippy: What's Happening to the Infamous Office Assistant in Office XP" [Introduced in Office '97]

<https://news.microsoft.com/2001/04/11/farewell-clippy-whats-happening-to-the-infamous-office-assistant-in-office-xp/>

REDMOND, Wash., April 11, 2001 — Whether you love him or you hate him, say farewell to Clippy automatically popping up on your screen.

Clippy is the little paperclip with the soulful eyes and the Groucho eyebrows. The electronic ham who politely offers hints for using Microsoft Office software.

But, after four years on-screen, Clippy will lose his starring role when Microsoft Office XP debuts on May 31. Clippy, the Office Assistant introduced in Office 97, has been demoted in Office XP. The wiry little assistant is turned off by default in Office XP, but diehard supporters can turn Clippy back on if they miss him.

Lisa Gurry, a Microsoft product manager explained: "Office XP is so easy to use that Clippy is no longer necessary, or useful. With new features like smart tags and Task Panes, Office XP enables people to get more out of the product than ever before. These new simplicity and ease-of-use improvements really make Clippy obsolete." she said.

"He's quite down in the dumps," Gurry joked. "He has even started his own campaign to try to get his old job back, or find a new one."

A report in USA Today in 2002 stated that Microsoft banked on its customers' contempt of Clippy to promote Office XP: "On Thursday, Microsoft is scheduled to unveil the last installment in a nontraditional advertising campaign that aims to sell the newest version of Office, called XP, by encouraging customers' hatred of Clippy." [Really, though... little encouragement was necessary.]

So, here we are, 18 years later and, wouldn't you know it, "Clippy's" lobbying to return to the limelight appears about to pay off.

"Clippy" is about to make a not-long-awaited comeback for Microsoft's Teams app. The effort is open source and on Github, so the animations are all publicly available. And I have to confess that Clippy is such a meme from the past that I would LOVE to be able to decorate the occasional iMessage with a bit of contemporary Clippy art. I hope someone ports those Clippy animations to iOS. :)



## Operation "ShadowHammer"

First, I'm going to share Kaspersky's post about this incident. Then I'll explain what puzzles me so much about this:

<https://securelist.com/operation-shadowhammer/89992/>

Earlier today, Motherboard published a story by Kim Zetter on Operation ShadowHammer, a newly discovered supply chain attack that leveraged ASUS Live Update software.

*"Hackers Hijacked ASUS Software Updates to Install Backdoors on Thousands of Computers"*  
The Taiwan-based tech giant ASUS is believed to have pushed the malware to hundreds of thousands of customers through its trusted automatic software update tool after attackers compromised the company's server and used it to push the malware to machines.

[https://motherboard.vice.com/en\\_us/article/pan9wn/hackers-hijacked-asus-software-updates-to-install-backdoors-on-thousands-of-computers](https://motherboard.vice.com/en_us/article/pan9wn/hackers-hijacked-asus-software-updates-to-install-backdoors-on-thousands-of-computers)

While the investigation is still in progress and full results and technical paper will be published during SAS 2019 conference in Singapore, we would like to share some important details about the attack.

In January 2019, we discovered a sophisticated supply chain attack involving the ASUS Live Update Utility. The attack took place between June and November 2018 and according to our telemetry, it affected a large number of users.

ASUS Live Update is an utility that is pre-installed on most ASUS computers and is used to automatically update certain components such as BIOS, UEFI, drivers and applications. According to Gartner, ASUS is the world's 5th-largest PC vendor by 2017 unit sales. This makes it an extremely attractive target for APT groups that might want to take advantage of their userbase.

Based on our statistics, over 57,000 Kaspersky users have downloaded and installed the backdoored version of ASUS Live Update at some point in time. We are not able to calculate the total count of affected users based only on our data; however, we estimate that the real scale of the problem is much bigger and is possibly affecting over a million users worldwide.

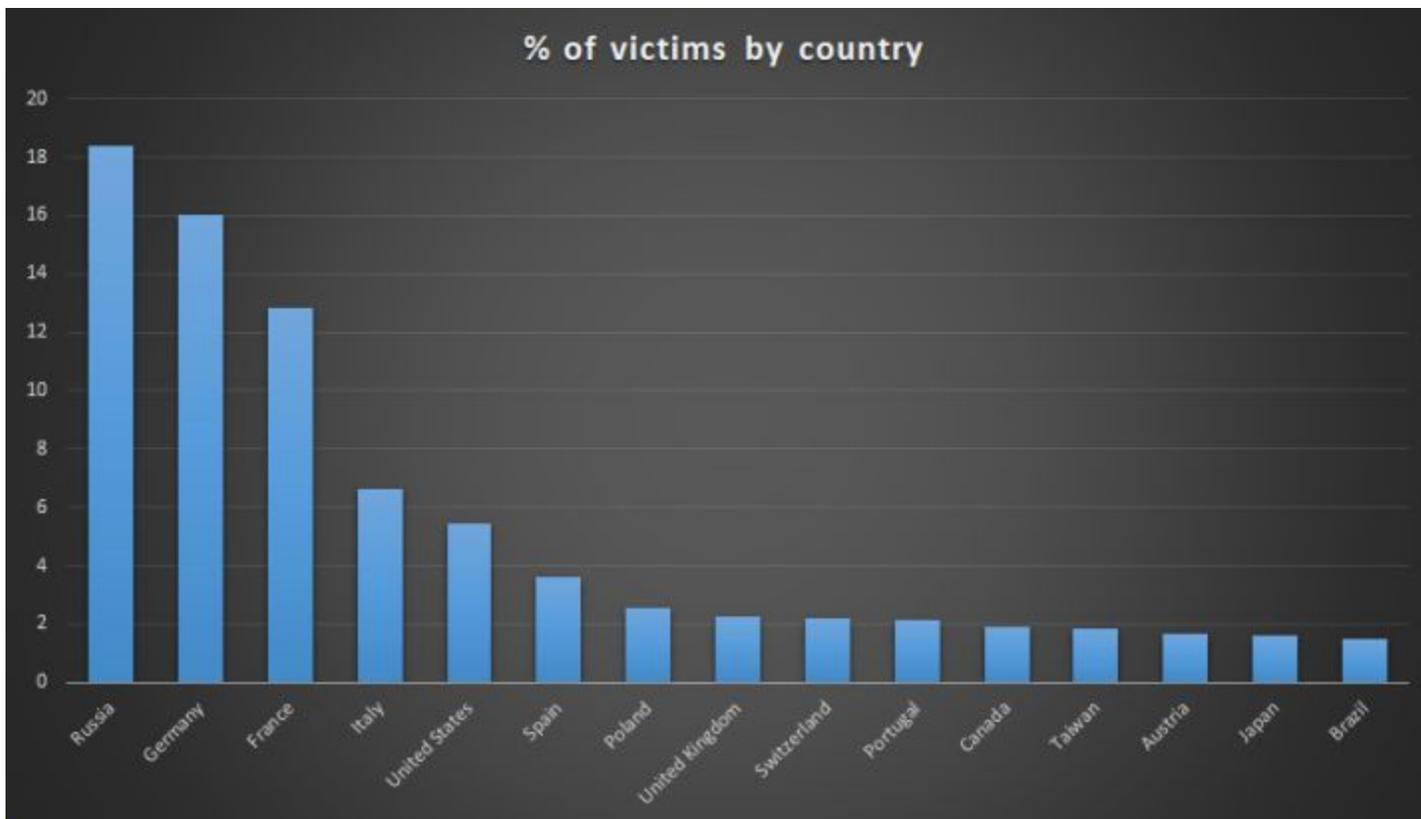
The goal of the attack was to surgically target an unknown pool of users, which were identified by their network adapters' MAC addresses. To achieve this, the attackers had hardcoded a list of MAC addresses in the trojanized samples and this list was used to identify the actual intended targets of this massive operation. We were able to extract more than 600 unique MAC addresses from over 200 samples used in this attack. Of course, there might be other samples out there with different MAC addresses in their list.

We believe this to be a very sophisticated supply chain attack, which matches or even surpasses the Shadowpad and the CCleaner incidents in complexity and techniques. The reason that it stayed undetected for so long is partly due to the fact that the trojanized updaters were signed with legitimate certificates (eg: "ASUSTeK Computer Inc."). The malicious updaters were hosted

on the official liveupdate01s.asus.com and liveupdate01.asus.com ASUS update servers.

Although precise attribution is not available at the moment, certain evidence we have collected allows us to link this attack to the ShadowPad incident from 2017. The actor behind the ShadowPad incident has been publicly identified by Microsoft in court documents as BARIUM. BARIUM is an APT actor known to be using the Winnti backdoor. Recently, our colleagues from ESET wrote about another supply chain attack in which BARIUM was also involved, that we believe is connected to this case as well.

A victim distribution by country for the compromised ASUS Live Updater looks as follows:



It should be noted that the numbers are also highly influenced by the distribution of Kaspersky users around the world. In principle, the distribution of victims should match the distribution of ASUS users around the world.

We've also created a tool which can be run to determine if your computer has been one of the surgically selected targets of this attack. To check this, it compares MAC addresses of all adapters to a list of predefined values hardcoded in the malware and alerts if a match was found.

<https://kas.pr/shadowhammer> (downloads a 50k ShadowHammerCheck.zip)

Also, you may check MAC addresses online ( <https://shadowhammer.kaspersky.com/> ). If you discover that you have been targeted by this operation, please e-mail us at: shadowhammer@kaspersky.com

Something feels a bit fishy about this.

Asus official servers were being used to supply the initial malware. And the malware was signed by legitimate ASUS certificates. And though Kaspersky's brief summary didn't mention it, other coverage noted that ASUS was being uncooperative in the extreme about this, denying that anything had happened.

But what I find so puzzling and curious is that the malware delivered by ASUS servers, from ASUS and signed by ASUS, used the victim's MAC addresses to identify individual specific ASUS machines. And what's most troubling is that no one but ASUS -- the manufacturer of those machines -- would reliably know what MAC addresses specific machines have.

As we know, MAC addresses are not highly secret, but neither are they widespread. A machine's MAC address is often printed on the label outside the box, and on the label of the machine itself. But MAC addresses are highly local because they do not transit the Internet. They provide local Ethernet network hardware addressing for use within a single Ethernet subnet. So, for example, any IP (Internet Packet) router serves as an intelligent "link" between separate Ethernet networks, with a different network on each of its interfaces. Unless the router is bridging two network at the Ethernet layer, the MAC address from one network is removed and its contained IP packet is routed to another interface, where it is encapsulated within an Ethernet packet with source and destination MAC addresses for THAT network.

My point is... How would some random external malicious agency obtain the physical hardware Ethernet MAC addresses for a large collection of specific ASUS machines?

One possibility is that these were wireless laptops. If so, then they would have been promiscuously broadcasting their MAC addresses more or less constantly to every WiFi access point within range. And remember that a MAC address is a 48-bit value composed of two halves, a 24-bit registered manufacturer number and a 24-bit serial number within that manufacturer. So the fact of them being ASUS laptops would have been evident from their MAC addresses. So there is some possibility that the machine MAC addresses of specific individuals could have somehow been gathered over time.

But if these were wired desktop machines it's difficult to come up with any theory to explain how a remote 3rd party could obtain those machines' Ethernet MAC addresses. And if some agency was close enough to a wired machine to obtain its MAC address it would have been in physical proximity, so probably already inside the machine... thus no need for any remotely downloaded malware shenanigans in the first place.

Occam's Razor suggests that when confronted with a lack of definitive evidence, the simplest explanation is likely to be the best. And distressing as that is, this suggests that the entire thing was likely a covert and deliberate campaign on the part of ASUS. Only ASUS has the certs to sign their update downloads. From Motherboard's reporting: *"The attackers used two different ASUS digital certificates to sign their malware. The first expired in mid-2018, so the attackers then switched to a second legitimate ASUS certificate to sign their malware after this."*

What sort of security are we to believe exists at ASUS if they were NOT a willing (or begrudging) collaborator in this drama? The "attackers" first sign their malware with ASUS' super-secret and

protected code signing certificate. Then those attackers place that ASUS-signed malware onto both of ASUS' software update servers. Then later, as that first certificate nears expiration, the "attackers" (in quotes) simply obtain ASUS' newly updated code signing certificate, resign their malware with that updated cert, and replace the soon-to-be-expired malware on both of ASUS software download servers with freshly signed new malware? Really? That's what we're to believe? And ASUS had no knowledge of any of this. The least that seems feasible is that a well-placed person on the inside arranged for all of this. Except for the Mac addresses. That would be a very different region within this very large company.

And as for those MAC addresses, only ASUS would have the sales records for those machines which indicate who owns which machines with which MAC addresses.

The fact that the follow-up malicious backdoor payload was later sourced from elsewhere gives ASUS some plausible deniability. And Kaspersky indicated that attribution was unavailable at the moment. Plus, it's very easy to plant a bit of misdirecting information, which would have been in ASUS' interest.

Motherboard wrote:

*Motherboard sent ASUS a list of the claims made by Kaspersky in three separate emails last Thursday but has not heard back from the company.*

*But the US-based security firm Symantec confirmed the Kaspersky findings on Friday after being asked by Motherboard to see if any of its customers also received the malicious download. The company is still investigating the matter but said in a phone call that at least 13,000 computers belonging to Symantec customers were infected with the malicious software update from ASUS last year.*

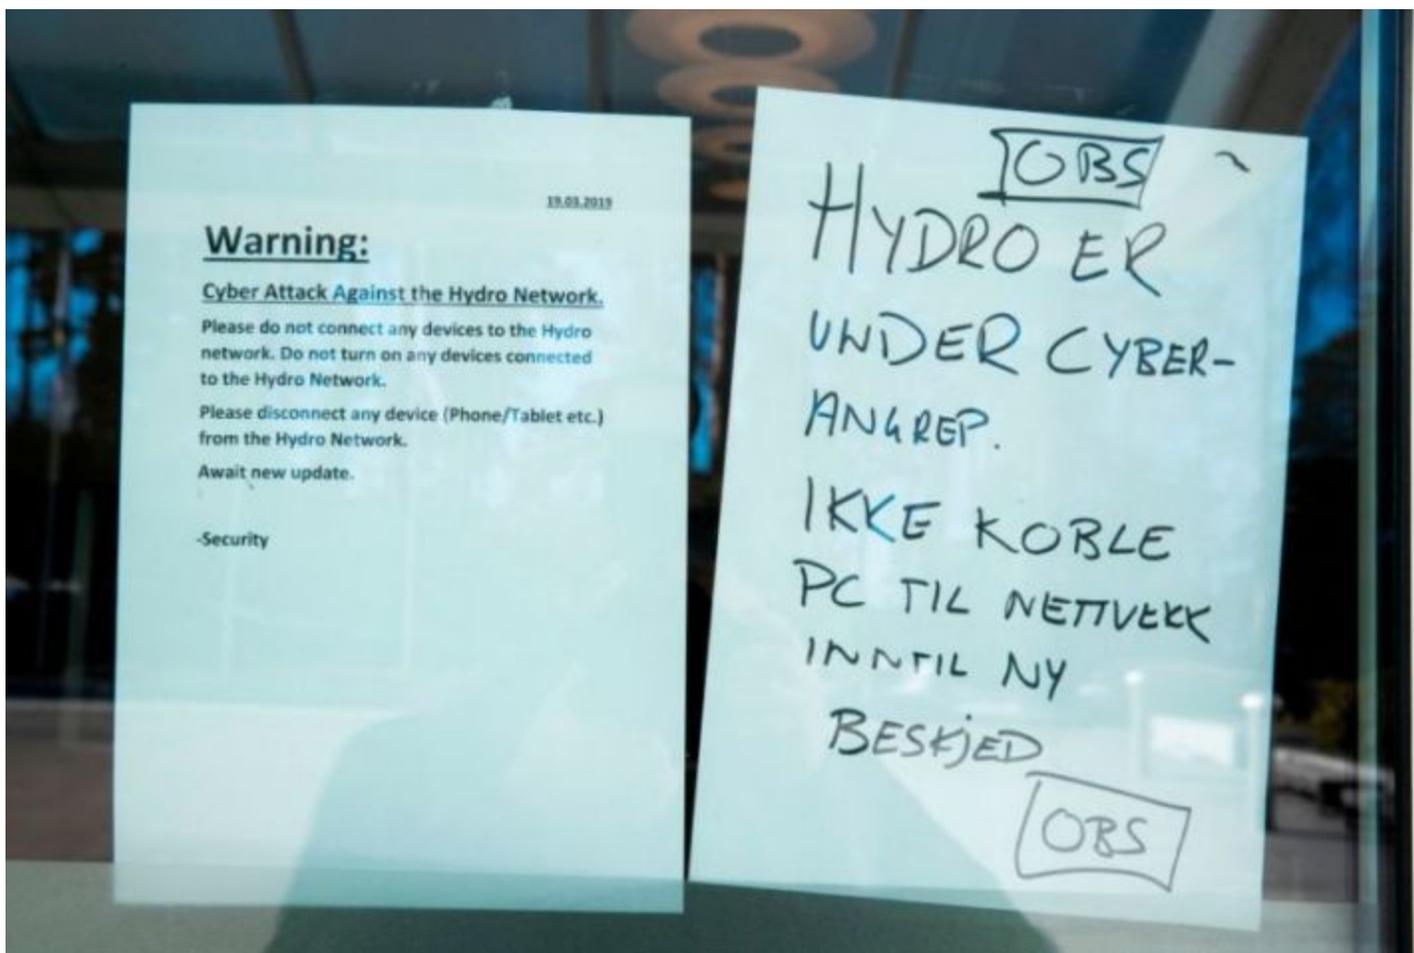
*Liam O'Murchu, director of development for the Security Technology and Response group at Symantec said: "We saw the updates come down from the Live Update ASUS server. They were trojanized, or malicious updates, and they were signed by ASUS."*

And, moreover, despite having been notified last January, Kaspersky noted that ASUS denied its servers were compromised when informed of the findings and continued to use one of the compromised certificates involved in the attack for at least a month after notification. And ASUS has still not invalidated the two compromised certificates, which means the attackers or anyone else with access to the un-expired certificate could still sign malicious files with it, and machines would view those files as legitimate ASUS files. (Except... it seems more likely that that neither of those certificates have been invalidated because they never left ASUS' control in the first place.

Perhaps ASUS is just huge and lumbering and too big to know better. But if so that's unsettling too.

Kaspersky apparently has additional details that we may learn in 12 days when they presents at their Security Analysts' Summit in Singapore.

## Norsk Hydro Ransomware Attack



<https://arstechnica.com/information-technology/2019/03/severe-ransomware-attack-cripples-big-aluminum-producer/>

Paraphrasing from ArsTechnica's coverage:

One of the world's biggest producers of aluminum has been hit by a serious ransomware attack that shut down its worldwide network, stopped or disrupted plants, and sent IT workers scrambling to return operations to normal.

Norsk Hydro of Norway said the malware first hit computers in the United States on Monday night. By Tuesday morning, the infection had spread to other parts of the company, which operates in 40 countries. Company officials responded by isolating plants to prevent further spreading. Some plants were temporarily stopped, while others, which had to be kept running continuously, were switched to manual mode when possible. The company's 35,000 employees were instructed to keep computers turned off but were allowed to use phones and tablets to check email.

Chief Financial Officer Eivind Kallevik said during a press conference Tuesday: "Let me be clear: the situation for Norsk Hydro is quite severe. The entire worldwide network is down, affecting our production as well as our office operations. We are working hard to contain and solve this situation and to ensure the safety and security of our employees. Our main priority now is to ensure safe operations and limit the operational and financial impact."

According to Kevin Beaumont, tweeting in his capacity as an independent researcher and citing local media reports, the ransomware that infected Norsk Hydro is known as LockerGoga. He said LockerGoga doesn't rely on the use of network traffic or on domain name system or command and control servers all which allow ransomware to bypass many network defenses.

An independent research group calling itself MalwareHunterTeam pointed to a LockerGoga sample uploaded to VirusTotal from Norway on Tuesday morning. At the time the malware was first scanned, it was detected by only 17 of the 67 biggest AV products, although detections increased once awareness of the Norsk Hydro infection grew. The malware had also once been digitally signed by security company Sectigo, but the certificate was revoked at an unknown time.

In a statement, Sectigo Senior Fellow Tim Callan wrote: "As a policy Sectigo revokes certificates used in malware attacks and does not issue certificates to known malware purveyors. We encourage security researchers to report instances of malware employing Sectigo certificates at [signedmalwarealert@sectigo.com](mailto:signedmalwarealert@sectigo.com)."

[When I first read this I thought to myself "Who the heck is "Sectigo"??? Guess who?! Our old friends "Comodo", now operating under a shiny new name... because they so thoroughly ruined their previous name.]

A text file that attackers included with the malware said:

There was a significant flaw in the security system of your company. You should be thankful that the flaw was exploited by serious people and not some rookies. They would have damaged all your data by mistake or for fun.

Your files are encrypted with the strongest military algorithms RSA4096 and AES-256. Without our special decoder it is impossible to restore that data. Attempts to restore your data with third-party software as Photorec, RannohDecryptor etc. will lead to irreversible destruction of your data.

The Norsk Hydro CFO said the majority of the company's plants were operating normally but that the network shutdown prevented plants from receiving future orders from customers. He said the losses at the moment were "minimal," but he conceded they would grow over time if automated systems aren't restored. Kallevik was unable to provide any timetable for how long it would take to disinfect the network.

He said company IT teams are working to remove the ransomware from infected systems. Once that's done, the teams plan to restore lost data using company backup systems, which Kallevik described as "good." Asked by a reporter if the company would rule out paying the demanded ransom, the CFO said the "main strategy is to use backup."

Lawrence Abrams at BleepingComputer, who is everyone's "go to" site for ransomware details added: "It should be noted that while this ransomware has had high profile targets, it is not the most active one out there targeting companies and has not seen wide distribution. Furthermore, it's very noisy as it consumes a lot of CPU, causes [Windows] explorer to crash repeatedly, and borks the system enough while encrypting that you can't run normal programs. Unless its launched on an idle machine, it would have a good chance of being spotted.

## **Microsoft renames “Windows Defender Advanced Threat Protection (ATP)” to... ... “Microsoft Defender Advanced Threat Protection (ATP)”**

<https://techcommunity.microsoft.com/t5/Windows-Defender-ATP/Announcing-Microsoft-Defender-ATP-for-Mac/ba-p/378010>

Posted last Thursday, Microsoft's blog posting is titled: “Announcing Microsoft Defender ATP for Mac” [That’s right... Microsoft A/V for macOS.]

Today, we’re announcing our advances in cross-platform next-generation protection and endpoint detection and response coverage with a new Microsoft solution for Mac. Core components of our unified endpoint security platform, including the new Threat & Vulnerability Management also announced today, will now be available for Mac devices.

We’ve been working closely with industry partners to enable Windows Defender Advanced Threat Protection (ATP) customers to protect their non-Windows devices while keeping a centralized “single pane of glass” experience. Now we are going a step further by adding our own solution to the options, starting with a limited preview today.

As we bring our unified security solution to other platforms, we’re also updating our name to reflect the breadth of this expanded coverage: Microsoft Defender ATP.

There are two key parts for cross-platform support for Microsoft Defender ATP on Mac:

A new user interface on Mac clients called Microsoft Defender ATP. The user interface brings a similar experience to what customers have today on Windows 10 devices.

Reporting for Mac devices on the Microsoft Defender ATP portal.

Microsoft Defender ATP can be installed on devices running macOS Mojave, High Sierra, or Sierra that you want to manage and protect.

In the limited preview, this app provides next-generation antimalware protection and allows end users to review and perform configuration of their protection, including: [[ All the standard A/V stuff we expect, quick and full scans, quarantining, etc. ]]

Users will also be able to configure advanced settings, for example:

- Disabling or enabling real-time protection, cloud-delivered protection, and automatic sample submission
- Adding exclusions for files and paths
- Managing notifications when threats are found
- Manually checking for security intelligence updates

The Microsoft AutoUpdate service is also installed, which ensures that the app is kept up-to-date and is properly connected to the cloud.

## Severe security bug found in most popular PHP library for creating PDF files

<https://www.zdnet.com/article/severe-security-bug-found-in-popular-php-library-for-creating-pdf-files/>

The three most popular libraries used by web servers to create PDF files for invoicing, purchase receipts, or whatever on the fly are TCPDF, mPDF and FPDF. And we now know, after the VERY responsible disclosure by a security researchers who waited not 6 hours, nor six days, nor six weeks... but six MONTHS after the flaw was fixed to disclose its existence... that a serious remote code execution flaw exists in versions of the most popular of the three PDF generators "TCPDF".

The vulnerability is a variation of another researcher's discovery: The flaw was first found by Sam Thomas, a researcher at Secarma who, in a series of experiments, showcased a deserialization bug affecting PHP apps over the summer of 2018. He released a research paper detailing PHP serialization attacks against the WordPress and Typo3 CMS platforms, but also the TCPDF library embedded inside the Contao CMS.

Then, in a blog post this past weekend <https://polict.net/blog/CVE-2018-17057> an Italian security researcher "Polict" revealed a new PHP serialization flaw impacting TCPDF in the same way as the one discovered by Thomas last summer.

Polict says the vulnerability he found can be exploited in two ways. The first case is on websites that allow user input to be part of the PDF file generation process, such as when adding names, eMail addresses or other details inside invoices which will be converted into a PDF.

The second is on websites that contain cross-site scripting (XSS) vulnerabilities where an attacker can plant malicious code inside the HTML source code that will be fed to the TCPDF library to convert into a PDF.

One way or another, the requirement is to supply specially malformed data to the TCPDF library which causes the TCPDF library to call the PHP server's "phar://" stream wrapper, and later abuse the PHP deserialization process to run code on the underlying server.

It is a complex attack routine requiring advanced PHP coding skills to exploit. And as we have discussed previously, deserialization exploits are difficult to uncover and are the bane of many programming languages, including Ruby, Java, and .NET, besides PHP. Recall that when we "serialize" we convert complex structured data into a linear, thus, serialized, byte stream. When we "deserialize" we are reversing the process to convert that byte stream back into the original complex data structure. But all too often the coders who write the deserializer assume that the byte stream they are being fed is valid and that it was created by a well-meaning serializer. And what is a deserializer??? It's an interpreter. It reads the incoming byte stream and interprets its meaning. As we have often noted here, interpreters are inherently fraught with problems.

So, anyway, Polict said that he reported the vulnerability (CVE-2018-17057) to the TCPDF library author last August. The TCPDF team then released an updated TCPDF v6.2.20 a month later September to address the issue. Unfortunately, when the TCPDF team did that they accidentally re-introduced the earlier vulnerability reported by Sam Thomas while attempting to patch the one reported by Polict. (Like I said... Interpreters are finicky beasts!)

So, finally, both issues were resolved in version 6.2.22.

Polict published details about this vulnerability just last Tuesday, after waiting a full six months after the patch's release due to the bug's severity and to allow websites and web app owners sufficient time to patch.

In their coverage of this, ZDNet notes that the TCPDF library is one of today's most popular PHP libraries and has been used *<quote> all over the place --in standalone websites, in content management systems (CMSs), CMS plugins, CMS themes, enterprise intranets, CRMs, HRMs, invoicing solutions, many PDF-centered web apps, and others. And patching isn't as easy as it sounds. In some cases, this might mean replacing a file and editing a build instruction, but in other places, this might require rewriting large swaths of code. </quote>*

As we know, unsexy PDF libraries like this are exactly the sorts of things that tend to be overlooked by web developers and website maintainers. So the takeaway for our listeners is: if you are responsible for any website that generates on-the-fly PDFs, make sure that if you are using TCPDF it's v6.2.22 or later.

### **Microsoft's Leaked Edge Browser is looking REALLY good**

<https://www.bleepingcomputer.com/news/microsoft/microsofts-leaked-edge-browser-should-make-google-worried/>

BleepingComputer reports that: Over the weekend, a leaked build for the Chromium-based Edge browser has been released that is providing users with their first look at the upcoming browser from Microsoft. If you are currently using Chrome, the reports indicate that this Edge preview browser feels, performs, and basically offers the same feature set.

Lawrence Abrams writes that: Microsoft has been quiet regarding their upcoming Microsoft Edge Insider browser, but a slow trickle of leaks has provided a bit more information. With this leaked build, though, users get their first full look at the upcoming Edge browser, which from all reports feels like it has the best chance of putting a dent in Google Chrome's market share.

Users who have tested the leaked build have also stated that browser performs really well when browsing the web and that it is more than ready for public preview. Microsoft has modified the layout of the browser to make it feel more like a Microsoft application. For example, the Settings pages have a left hand navigation bar similar to other Windows 10 apps. Microsoft also included their own services into the browser. For example, Google Safe Browsing has been removed in favor of Microsoft's SmartScreen.

#### Edge supports Chrome Extensions

In addition to setting up a dedicated Microsoft Extension store, Edge also allows users to enable the installation of extensions from Chrome's web store. While they state that these extensions are unverified as a warning, it provides an enormous pool of extensions for users to install right from release.

This feels significant. The largest browser share and the browser built into the largest desktop and laptop share effectively merge to being one.

## **From our "pained abbreviation" department comes: SHAKEN/STIR**

SHAKEN stands for: Signature-based Handling of Asserted Information Using toKENs (SHAKEN).  
STIR comes from: Secure Telephone Identity Revisited.

Horrible as those abbreviations are, together SHAKEN and STIR deliver a protocol for authenticating phone calls with the help of cryptographic certificates.

The US Federal Communications Commission has been pushing for SHAKEN/STIR's adoption and has imposed the end of 2019 as a hard deadline for networks implementing the protocol.

<https://www.fcc.gov/call-authentication>

"Combating Spoofed Robocalls with Caller ID Authentication"

FCC Chairman Ajit Pai: "American consumers are sick and tired of unwanted robocalls, this consumer among them. Caller ID authentication will be a significant step towards ending the scourge of spoofed robocalls. It's time for carriers to implement robust caller ID authentication."

How Will Caller ID Authentication Help Consumers?

Caller ID authentication is a new system aimed at combating illegal caller ID spoofing. Such a system is critical to protecting Americans from scam spoofed robocalls and would erode the ability of callers to illegally spoof a caller ID, which scam artists use to trick Americans into answering their phones when they shouldn't. Additionally, consumers and law enforcement alike could more readily identify the source of illegal robocalls and reduce their impact. Industry stakeholders are working to implement caller ID authentication, which is sometimes called SHAKEN/STIR. Once implemented, it should greatly help the accuracy of caller ID information and should provide consumers with helpful information for determining which calls are authenticated.

SHAKEN/STIR is a framework of interconnected standards. SHAKEN/STIR are acronyms for Signature-based Handling of Asserted Information Using toKENs (SHAKEN) and the Secure Telephone Identity Revisited (STIR) standards. This means that calls traveling through interconnected phone networks would have their caller ID "signed" as legitimate by originating carriers and validated by other carriers before reaching consumers. SHAKEN/STIR digitally validates the handoff of phone calls passing through the complex web of networks, allowing the phone company of the consumer receiving the call to verify that a call is from the person making it.

In November 2018, Chairman Pai sent letters to voice providers asking those that apparently had not yet established concrete plans to protect their customers using the SHAKEN/STIR standards to do so without delay. In February 2019, Chairman Pai welcomed many carriers' commitment to meeting his timeline for implementation, called on others to "catch up," and made clear that the FCC would consider regulatory intervention if necessary.

Last Wednesday, AT&T and Comcast announced that they had successfully tested what they believe to be the first SHAKEN/STIR-authenticated call between two different telecom networks.

Work on the SHAKEN/STIR protocol has been underway for a while, and telecom operators have tested it internally, but only for calls originated and terminated within their own networks where they could verify that everything was working as intended and calls were getting signed correctly when made, and verified correctly when received.

Once broadly adopted, incoming calls NOT signed can be flagged as such and at some point, when all valid calls are being signed, unsigned calls can simply be dropped as a subscriber option. And spoofed caller ID for signed calls will be impossible. This will not, by itself, stop all RoboCalling... But it will chill the callers significantly to know that law enforcement

## **PuTTY Releases Important Software Update to Patch 8 High-Severity Flaws**

Putty 0.71 Fixes Weakness That Allows Fake Login Prompts

The popular SSH client program PuTTY has released an updated version of its software that includes security patches for 8 high-severity vulnerabilities.

As many of us know, PuTTY is one of the most popular and widely used open-source SSH clients for secure network access of remote computers over SSH, and with less security over Telnet.

Nearly two years after releasing the previous version of its software, on March 16th the developers of PuTTY released version 0.71 for Windows and Unix operating systems. According to an advisory available on their website, all previous versions of the PuTTY software have been found vulnerable to multiple security vulnerabilities that could allow a malicious server or compromised server to hijack the PuTTY client's system in many different ways.

I won't go into the details here, The hacker News has very good coverage of them for anyone who is interested: <https://thehackernews.com/2019/03/putty-software-hacking.html>

The PuTTY release v0.71 page:

<https://www.chiark.greenend.org.uk/~sgtatham/putty/latest.html>

The takeaway for our listeners is that while the risks are low, if you are a user of PuTTY you'll likely be glad to know that a handful of known problems have been found and fixed and that you can easily and freely update your current installed version.

## **Firefox to Import Windows Root Certs To Avoid Antivirus SSL Scanning Issues**

<https://www.bleepingcomputer.com/news/software/firefox-to-import-windows-root-certs-to-avoid-antivirus-ssl-scanning-issues/>

When we were exploring the recent mess caused by the Firefox update and the AV vendors who had added their A/V HTTP TLS interception roots into the Windows root store... but not to Firefox's, we noted that Firefox DID have an option to import the Windows root store: `security.enterprise_roots.enabled` ... set to: true

Now Mozilla is considering automatically flipping that preference in detectable situations to avoid future problems with 3rd-party A/V...

[https://bugzilla.mozilla.org/show\\_bug.cgi?id=1533397](https://bugzilla.mozilla.org/show_bug.cgi?id=1533397)

"Retention/Engagement impact of enabling the Enterprise roots feature in presence of an AV"

Several AVs recently broke HTTPS with their HTTPS scanning features that require their certs to be added to our cert store (Avast on bug 1523701, BitDefender on bug 1508624, Kaspersky on bug 1449115). The security team confirmed that having the preference `security.enterprise_roots.enabled` set to true would have fixed all of these issues without known regressions and we want to validate that in the presence of an AV, enabling this preference would have a positive impact on retention and engagement (we cannot detect a change in certificate error page displays through telemetry since telemetry is sent over https, that breaks in these instances....).

Description of the impacted population:

- Win 10 and Win 8 release users (the API allowing detection of an AV registered with the system was only available since Win8)
- An AV is registered with the system and IS NOT Windows Defender. This information is available on telemetry under "sec.antivirus"

Test cohort: Win 8+ users on release 66 with an AV registered that is not Windows Defender and don't have `security.enterprise_roots.enabled` set to true. Set `security.enterprise_roots.enabled` to true.

## **Windows 7 Gets an Extra Life With Windows Virtual Desktop**

<https://www.bleepingcomputer.com/news/microsoft/windows-7-gets-an-extra-life-with-windows-virtual-desktop/>

Last Thursday Microsoft announced that their new Windows Virtual Desktop product was now available for public preview. This technology allows the enterprise to move their desktops and applications into Azure for hosting on Windows 10 and Windows 7 Virtual Machine operating systems that are always secured with the latest updates.

Windows Virtual Desktop was being privately tested by select organizations while Microsoft was fine tuned the offering. Now, anyone with an Azure subscription, which you can get for free for a year, can test deploying apps and desktops to the cloud.

To use the service, administrators first create a tenant in Azure then publish desktops and applications to it. These available services will then be made available in "feeds" that users can subscribe to using a new Windows Virtual Desktop client called "Remote Desktop".

No one can understand why Microsoft would deliberately have the name collision between their Azure cloud offering and the existing Remote Desktop RDP client. Perhaps the plan is for a single "Remote Desktop" client to be able to connect to either an Azure cloud Windows instance, or to any other network-accessible machine over RDP.

And, for the many (very many) organizations that require Windows 7 and that are not yet willing to be forced to make the jump to Windows 10, Microsoft is offering an interesting incentive to use their Windows Virtual Desktop product: The Azure Cloud version of Windows 7 will automatically receive security updates at least until 2023... So for an additional three years.

This would allow organizations to upgrade their local workstations to Windows 10, while still continuing to utilize any required Windows 7 applications through a Windows Virtual Desktop.

### **KB4493132 Update Notifies Windows 7 Users of End of Support Date**

<https://www.bleepingcomputer.com/news/microsoft/kb4493132-update-notifies-windows-7-users-of-end-of-support-date/>

KB4493132 is the Windows Update for Windows 7 that will be notifying Windows 7 Users of the approaching End of Support Date for their operating system versions.

It appeared last week and so far it is NOT selected for installation by default. But neither does it identify itself clearly. It states that it is an update to "resolve issues in Windows" rather than one that just displays "End of Life" notifications. So hopefully, unless and until it is selected and deliberately installed, it won't be added to any Win 7 system that doesn't want it.

If it SHOULD get into a system by mistake, Microsoft appears to have learned from their previous GWX (Get Windows 10) fiasco and they have included a very small fine print:

**[ x ] Do not remind me again** checkbox:



The notification window contains the message:

After 10 years, support for Windows 7 is nearing the end.

January 14, 2020 is the last day Microsoft will offer security updates and technical support for computers running Windows 7. We know can be difficult, that's why we're reaching out early to help you backup your files and prepare for what's next.

Lawrence Abrams has a detailed analysis of the download and executable and Windows scheduler and where everything lives. So I'm certain that it can be removed and/or neutered if it ever becomes a nuisance.

### **Sci-Fi: Amazon Prime: *The Expanse***

Ryk Brown's "**The Frontiers Saga**" series. Now on book #26.  
"Comfort Food" -- this is "comfort reading."

### **SQRL**

There had been discussion in the SQRL web forums about remaining logged into a site. That's not in any way related to SQRL, but the easy of sign in that SQRL provides does sort of shift the balance away from remain signed-in everywhere all the time since it's so easy now to reassert one's identity. Anyway, in this discussion I saw two posts that made me smile:

1. Day before yesterday at 1pm: "Gristle" wrote: I am loving logging in with SQRL so much, that I actually don't want the cookie to persist longer than an hour, or maybe even a few minutes. It's just so fun and simple to use SQRL!
2. Yesterday he quoted himself and added: I'm not joking, I find myself logging out and in and out and in just for fun. It's crazy, I know, but it's just so cool. Reminds me of the first time I tried unlocking a phone with my fingerprint. I kept locking and unlocking it, unbelieving that there is so much crypto behind such a simple gesture.

### **SpinRite**

*I have a few questions about SpinRite.*

*Those sectors you read with surface scan errors -- do you try to overwrite them so that sector mapping by S.M.A.R.T takes place? Recall the remapping will only take place when the data is overwritten which in a file system may never happen until the disk is completely (sector-by-sector) reformatted -- imagine a file that is read often but never re-written and resides on one or more sectors with surface scan errors.*

If you think about this a bit you'll see how this can work:

Surface defects do not manifest when sectors are written because "writing" to the disk is an entirely blind process. The drive gets no information about the readability of a freshly-written sector. That only occurs when a subsequent attempt is made to read-back that sector. Because a "read-after-write" would be prohibitively slow, all modern hard drives incorporate Error

Correction Code (ECC) technology which creates a static safety margin to defend against the possibility that, due to some problem which may have occurred during writing, or a surface anomaly that interferes with reading, the data that was written to a sector cannot be read back without the aid of algorithmic correction.

ECC operates by appending additional carefully designed redundant data to the end of the sector, based upon the sector's intended content. This extra data allows the location of a read-back problem to be identified along with its length and its error mask. From this data the drive is able to determine -- within limits -- what data was originally written to the sector. Those error correction recovery limits are the number of separate errors occurring within a single sector and each error's bit-run-length.

On today's drives, data densities have grown so high that some level of "background" error correction is occurring more or less continuously. It's only when the extent of correction required for read-back recovery begins to approach the drive's inherent limits that the drive gets worried and decides to remove that worrisome sector from future service.

*If you do overwrite them and force the mapping, how do you stop a disk from being rendered useless because the max limit (say 100) of remapping has taken place due to your reading all sectors and forcing all those remappings?*

A sector can be misread due to its own aging or transient vibration or physical shock which forced the head off-track when the sector was last written or when it is now read. To determine whether the error was transient and not caused by an actual defect, SpinRite first writes inverted data then reads it back. It then re-writes the original data and reads that back. If any of those three reads -- the original read, the inverted read, the re-inverted read -- result in sufficiently threatening data read-back errors, the drive will take that sector out of service and remap the sector.

When a sector is found with a surface scan error, are you able to tell me what NTFS (Windows, in my case) file is using that bad sector?

Since a bad sector is replaced with a good sector at the drive's physical interface level, this falls beneath the operating system's file system. In essence, all surface errors and remappings are transparent to the OS and file system. So nothing needs to be done there. Before this autonomous drive-level sector-replacement technology existed universally as it does now, SpinRite did this itself. All early versions of SpinRite managed surface defects for the drive. SpinRite would determine which cluster the defective sector occupied. It would then determine what role that cluster was playing in the file system, if any. If the cluster was in use by a directory or file, it would obtain the closest nearby free cluster, copy the bad cluster's data into the new cluster, and re-link the replacement cluster into the system's file system. It would mark the defective sector as bad in the low-level format and mark the defective cluster containing that bad sector as bad in the system's cluster management tables.

All of that technology still exists in today's SpinRite for possible use if it should encounter the need. But since SpinRite is able to interact with all modern drives to perform these data relocations underneath the file system, other than on floppy disks, SpinRite's high-level defect management is never needed today.

# Pwn2Own Vancouver 2019 - Tesla, Pwned

Organized by Trend Micro's ZDI (zero-day initiative) with support from Microsoft, Tesla and VMware, this year's Pwn2Own was held last Wednesday, Thursday and Friday in Vancouver. **The first day** resulted in 4 successful hacks and one partial win, with the contestants earning \$240,000 USD in cash awards – plus the laptops used to demonstrate their research. In their traditional blow-by-blow style, Trend Micro described the first day as follows:

*The contest started with the team of Fluoroacetate (Amat Cama and Richard Zhu) targeting the Apple Safari web browser. They successfully exploited the browser and escaped the sandbox by using an integer overflow in the browser and a heap overflow to escape the sandbox. The attempt nearly took the entire allowed time because they used a brute force technique during the sandbox escape. The code would fail then try again until it succeeded. The demonstration earned them \$55,000 USD and 5 points towards Master of Pwn.*

*The Fluoroacetate duo returned targeting Oracle VirtualBox in the virtualization category. Although their first attempt failed, the second attempt successfully used an integer underflow and a race condition to escalate from the virtual client to pop calc at medium integrity. It wasn't the race condition that caused their failed first attempt. Their memory leak was working, but their code execution failed. Everything aligned on the second attempt, which earned them \$35,000 USD and 3 more Master of Pwn points.*

*Next up, Pwn2Own newcomer anhdaden of STAR Labs also targeted Oracle VirtualBox. He also used an integer underflow to escalate from the virtual client to execute his code on the hypervisor at medium integrity. Interestingly, he used a unique integer underflow different than the previously demonstrated underflow. His first foray into Pwn2own netted him \$35,000 USD and 3 Master of Pwn points. This also marks the first Vietnamese winner at Pwn2Own. We hope to see more from him in the future.*

*In their final entry for Day One, the Fluoroacetate duo targeted the VMware Workstation. They leveraged a race condition leading to an Out-Of-Bounds write to go from the virtual client to executing code on the underlying host operating system. They earned \$70,000 USD and 7 additional Master of Pwn points. This brings their Day One total to \$160,000 and 15 Master of Pwn points.*

[Note that this is a full VMWare escape to execute code on the host OS, for which, as we recently noted here, Zerodium has upped their bounty to half a million dollars. So it seems clear that these guys are good guys who are interested in increasing security rather than indirectly attacking others.]

*The final entry in Day One saw the phoenix & qwerty team (@\_niklasb @qwertyoruiopz and @bkth\_) targeting Apple Safari with a kernel elevation. They demonstrated a complete system compromise. By browsing to their website, they triggered a JIT bug followed by a heap out-of-bounds (OOB) read – used twice – then pivoted from root to kernel via a Time-of-Check-Time-of-Use (TOCTOU) bug. Unfortunately, it was only a partial win since Apple already know of one of the bugs used in the demo. Still, they earned themselves \$45,000 USD and 4 points towards Master of Pwn.*

## **Day Two:**

*Our day began with the Fluoroacetate duo of Amat Cama and Richard Zhu targeting the Mozilla Firefox web browser. They leveraged a JIT bug in the browser, then used an out-of-bounds write in the Windows kernel to effectively take over the system. They were able to execute code at SYSTEM level just by using Firefox to visit their specially crafted website. The effort earned them another \$50,000 and five more points towards Master of Pwn.*

*The prolific duo returned with perhaps their greatest challenge of the competition. Starting from within a VMware Workstation client, they opened Microsoft Edge and browsed to their specially crafted web page. That's all it took to go from a browser in a virtual machine client to executing code on the underlying hypervisor. [Do you realize how much that would have meant to Zerodium? They would have moved heaven and earth to obtain that hack. Anyway...] They started with a type confusion bug in the Microsoft Edge browser, then used a race condition in the Windows kernel followed by an out-of-bounds write in VMware workstation. The masterfully crafted exploit chain earned them \$130,000 and 13 Master of Pwn points. They now have a commanding lead with 33 points total. In the two days of the competition, they have racked up a total of \$340,000 as a result of their phenomenal work. Tomorrow, they will attempt to cap their week off with a successful demonstration in the automotive category.*

*The third attempt of the day had Niklas Baumstark (@\_niklasb) target the Mozilla Firefox web browser. He used a JIT bug in the browser followed by a logic bug to escape the sandbox. In a real-world scenario, an attacker could use this to run their code on a target system at the level of the logged-on user. The successful demonstration earned him \$40,000 and 4 Master of Pwn points.*

*The final attempt for Day Two had Arthur Gerkis (@ax330d) of Exodus Intelligence targeting Microsoft Edge. Another newcomer to Pwn2Own, he wasted no time by using a double free bug in the renderer followed by a logic bug to bypass the sandbox. His debut entry earned him \$50,000 and five points towards Master of Pwn.*

*That brings Day Two to a close. We awarded \$270,000 for 9 unique bugs today, which brings the two-day total to \$510,000. Join us tomorrow as we debut the automotive category with the two final entries of Pwn2Own Vancouver 2019.*

## **Day Three:** The Automotive Category.

Two months ago, writing for Forbes Magazine, Thomas Brewster wrote:

*Think you can hack a Tesla? Now's your chance. And you could win more than \$900,000 in the process.*

*For the first time ever, Pwn2Own, perhaps the world's best-known competition for ethical hackers, will have a Tesla Model 3 opened up for participants to break. Prizes range from \$35,000 to \$250,000. The more difficult the hack, the greater the prize.*

*The lowest prize will go to an as-yet-unspecified attack on the electric car's infotainment*

*system. The top \$250,000 reward will go to the first person or team who can break any of three critical Tesla internals: the Gateway, the Autopilot or the VCSEC.*

*The Gateway acts as the central hub for controlling data flowing around the Tesla. Taking control of that system would give a hacker power over many of the car's functions. Manipulating Tesla Autopilot could lead to all too obvious problems. Imagine if the hacker simply shut Autopilot down without the driver noticing.*

*The VCSEC is the part of a Tesla responsible for a variety of security functions, including the alarm. Again, it's not hard to guess just what a hacker could do if they commandeered that part of the car.*

*Another \$100,000 is on offer for the first to hack the doors off by breaking the key fob or mobile app unlock tech. Starting the car without owning the legitimate key will also land a lucky hacker \$100,000.*

So what happened on day three???

*[ZDI writes], the day began not with a bang, but with a whimper as the Team KunnaPwn withdrew their entry from the automotive category. Although they didn't demonstrate any of their research at this contest, we hope they submit some of their research to our program in the future.*

[Note: This was supposed to be the day when Team KunnaPwn demonstrated a hack of the Tesla Model 3's "VCSEC" security component, but they withdrew from the competition. "VCSEC" stands for Vehicle Controller Secondary and is responsible for security functions such as the alarm.]

However, Fluoroacetate was up next...

*When their scheduled time arrived, the dynamic Fluoroacetate duo of Richard Zhu and Amat Cama thrilled the assembled crowd as they entered the vehicle. After a few minutes of setup, and with many cameras rolling, they successfully demonstrated their research on the Tesla Model 3 [Chromium based] Internet browser. They used a JIT bug in the renderer to display their message [on the Tesla's infotainment system] and earn \$35,000. Of course, this is Pwn2Own so they also get the car.*

-----

*Overall, the three days of Pwn2Own Vancouver 2019 have been a great success. We have awarded a total of \$545,000 for 19 unique bugs in Apple Safari, Microsoft Edge and Windows, VMware Workstation, Mozilla Firefox, and – in its inaugural year – the Tesla infotainment system.*

*And it should come as no surprise that the Fluoroacetate team of Richard Zhu and Amat Cama have been crowned the Master of Pwn for 2019! Their amazing research earned them \$375,000 over the contest and resulted in 36 Master of Pwn points. They dominated Pwn2Own Tokyo and have carried that wave through to the spring. We can't wait to see what's next for this pair of talented researchers.*