Security Now! #710 - 04-16-19
DragonBlood

This week on Security Now!

This week we discuss a malicious use of the URL tracking "ping" attribute, more on WinRAR, more 3rd-party A/V troubles with Microsoft, and other new trouble from last week's patch Tuesday, also the good things that patch Tuesday accomplished for Microsoft and for Adobe, another security-tightening change being proposed by Google, Russia's Roskomnadzor finally lowers the boom on Facebook and the incredible Taj Mahal APT framework. We then touch on a bit of miscellany, answer a SpinRite upgrade question, share some closing the loop feedback from our listeners... and then we look at DragonBlood, the first effective attack on the new WPA3 protocol (which didn't take long.)

Windows 10 Start Menu immediately after installing Win10 “Professional”

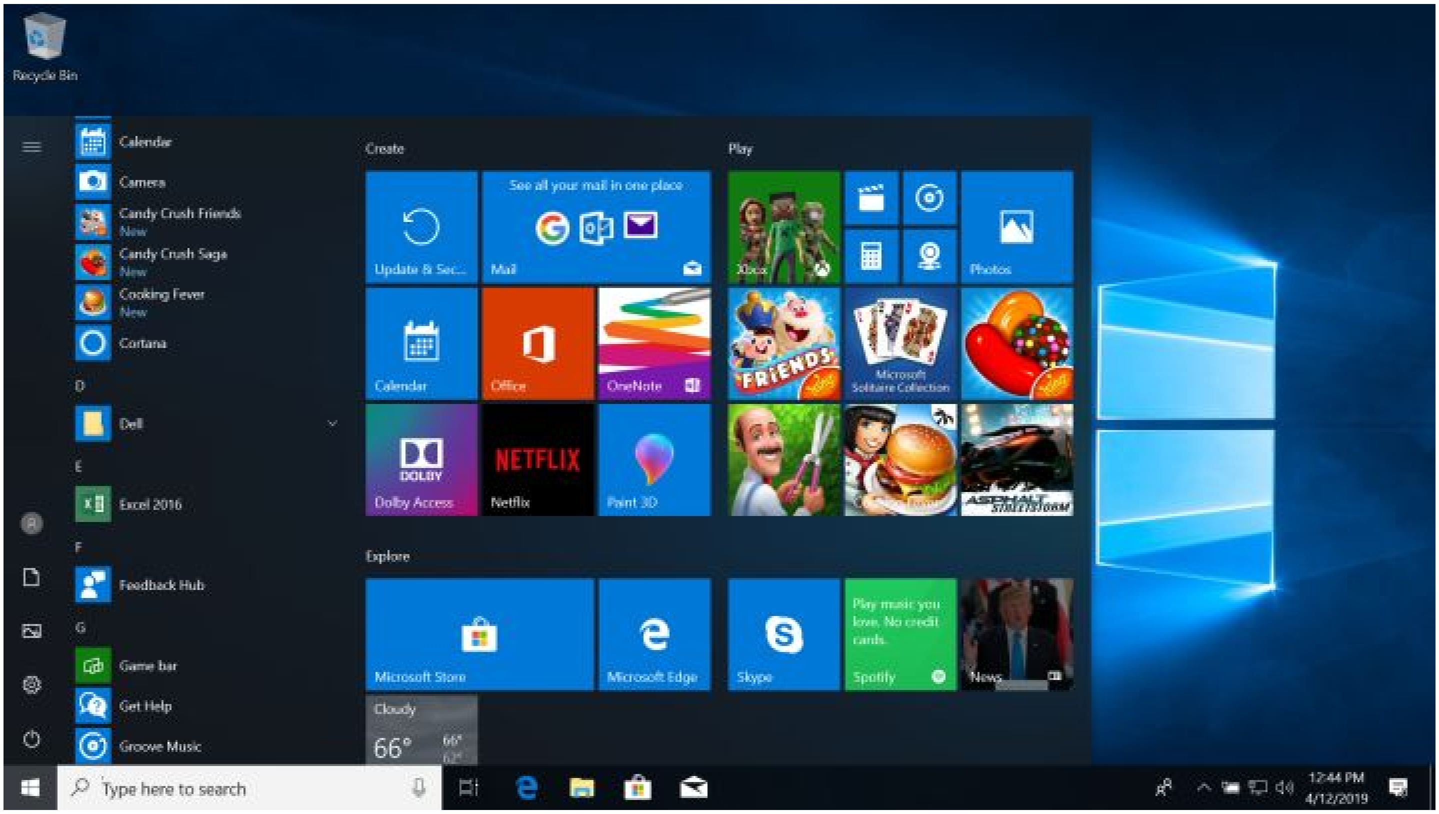

Win10 Start Menu after installing Win10 Long Term Service Channel (LTSC)

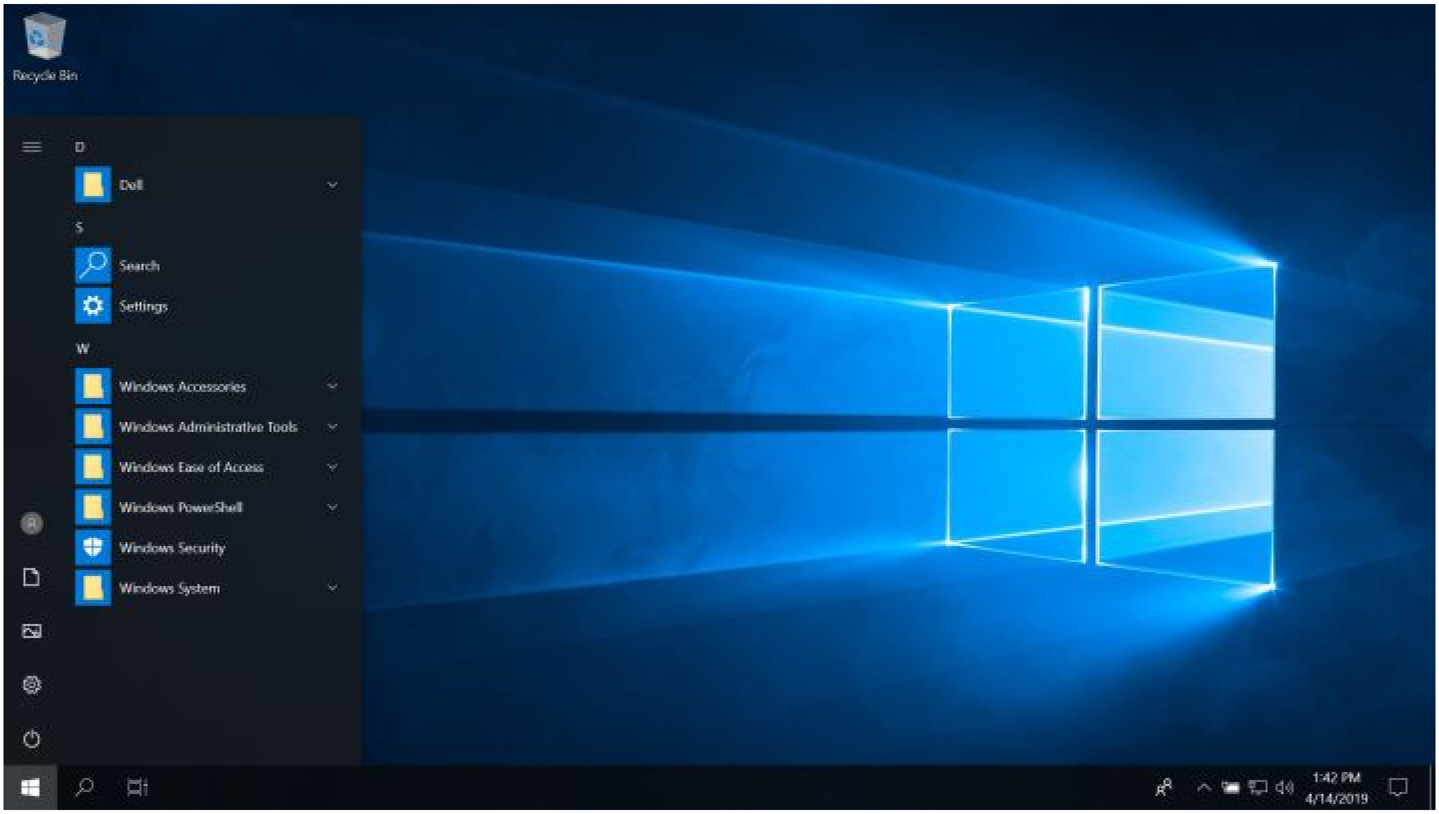

Security News

So I need to start right out by acknowledging a failure of my imagination.
Our listeners will recall that last week's podcast topic was "URL "Ping" Tracking" where I described the HTML5 feature of the "ping" term into <a> anchor tags with URLs for the sole purpose of tracking a user's hyperlink clicks rather than having the link's URL take them on a redirection journey.

My imagination failed me because I stated that the SOLE PURPOSE for the "ping" term was sending an asynchronous POST query to... Wherever.

And while that WAS its sole intended purpose, what never occurred to me was that this simple and benign mechanism -- albeit annoying from an anti-tracking privacy standpoint -- could be weaponized to generate DDoS attacks! Last week I noted that it deliberately ignores and same-origin policy enforcement and allows the "pings" to be sent anywhere.

This fact allows Javascript to edit the PING URL and to programmatically click the URL to launch a PING query at ANY other website.

Imperva research has uncovered a DDoS attack which utilized these HTML pings to perform distributed denial of services attacks on various sites. In one attack, which peaked at 7500 requests per second, a total of 70 million requests were generated from approximately 4,000 IP address over the course of 4 hours.

Since Safari and Opera offer no provision for disabling this behavior, and since it's enabled by default in Chrome and Google plans to remove the provision to disable it (not that that would matter, since no users would do so), it looks like our browser designers will need to come up with a way to preserve this functionality while preventing its abuse.

The WinRAR Nightmare

As Sophos Naked Security site puts it: "An ancient WinRAR vulnerability made public in February is now well on its way to becoming one of the most widely and rapidly-exploited security flaws of recent times."

Their coverage from yesterday is titled: "Flood of exploits targetting ancient WinRAR flaw continues"

The latest evidence is a report from Microsoft’s Office 365 Threat Research team which identified it as being used by the ‘MuddyWater’ APT group to target organizations in the satellite and communications industry. WinRAR was far too tempting for cybercriminals to ignore, within days stirring up a hornet’s nest of exploits to the tune of 100 or more. Microsoft’s blog about recent targeted attacks serves as yet another warning to organisations or individuals who still haven’t updated -- or removed -- WinRAR yet. We need word to get around as much as possible since, as we've noted previously, unregistered users, or users who no longer maintain their registered eMail accounts, will have no way of being informed of this ongoing danger.

Microsoft detected the threat to Office 365 in early March. The ATP attackers use a Word attachment claiming to be from the Ministry of Foreign Affairs (MFA) of the Islamic Republic of Afghanistan as the lure. Opening this triggers a download from a OneDrive link (which has since been shutdown) to an archive containing a second Word file within which is embedded a macro initiating the payload in the form of a PowerShell script which opens a command backdoor allowing the attackers to deliver the malicious ACE file containing the exploit.

It’s a bit convoluted because the attackers need to induce the user via a bogus warning dialogue to restarting the PC for the full attack to get setup and running. And while the entire exploit chain won't succeed every time, it's a numbers game targeting multiple individuals inside a targeted organization. So it only takes one to get in. (A similar strategy certainly succeeded against Sony.)

Sophos notes:

No one should assume that just because the attacks detected so far have been connected to nation state actors this will always be the case. Commercial exploits won’t be far behind – WinRAR’s half a billion reported users is a lot of victims to aim at.

More 3rd-party A/V vs Microsoft problems...
Applying this month's patch Tuesday patches one week ago has resulted in widespread problems for the users of a number of major 3rd-party antivirus systems.

These widespread problems are causing Windows 7, Windows 8.1, Windows Server 2012, and Windows Server 2012 R2 to freeze, be unable to boot, or hang on installing updates. And it also appears that some Windows 10 users have also been affected.

According to support articles from Microsoft, Avast, Avira, and Sophos, there is a conflict between some of the recent updates and A/V software such as Sophos Endpoint Protection, ArcaBit, AVG Business Edition, Avira antivirus, and Avast for Business and CloudCare.

In Sophos' support article they state that the updates could cause Windows to fail to boot. Reports from users also indicate that the update process may hang at the Configuring Updates stage.

All of the actors are aware of the problems and at least in the case of Sophos, Microsoft has written: "To prevent further issues, Microsoft has placed a block on the conflicting updates so that they are not offered to users running Sophos Endpoint until a solution is made available."

Sophos is reporting conflicts with the following updates:

- April 9, 2019 - KB4493467 (Security-only update) - Windows 8.1, Windows Server 2012 R2
- April 9, 2019 - KB4493446 (Monthly Rollup) - Windows 8.1, Windows Server 2012 R2
- April 9, 2019 - KB4493448 (Security-only update) - Windows 7 Service Pack 1, Windows Server 2008 R2 Service Pack 1
- April 9, 2019 - KB4493472 (Monthly Rollup) - Windows 7 Service Pack 1, Windows Server 2008 R2 Service Pack 1
- April 9, 2019 - KB4493450 (Security-only update) - Windows Server 2012, Windows Embedded 8 Standard
- April 9, 2019 - KB4493451 (Monthly Rollup) - Windows Server 2012, Windows Embedded 8 Standard

Avast has indicated that the KB4493472 and KB4493448 updates are causing problems, as well as an additional KB4493435 update (which is a cumulative security update for Internet Explorer.)

Avira has stated that KB4493509 for Windows 10 and KB4493472 and KB4493448 for Windows 7 will cause Windows to operate slower than normal. They suggested uninstalling these Windows updates for now until they can release a fix.

How to resolve freezes and boot loops

If Windows fails to start, freezes, or gets stuck at Configuring Updates after installing the April 2019 updates, you may have to remove the update to get Windows working properly.

Sophos recommends the following steps to remove the updates from an affected computer:

- Boot into safe mode
- Disable the Sophos Anti-Virus service
- Boot into normal mode
- Uninstall the Windows KB
- Enable the Sophos Anti-Virus service
- If enabled, Tamper Protection will need to be disabled to re-enable the service

Avast recommends that you boot your computer into Safe Mode and then uninstall the KB4493472, KB4493448, and KB4493435 updates.

Günter Born of Borncity.com has also stated that there are reports that Windows 10 may be affected as well. He has had reports from Windows 10 users who are having similar issues when Sophos is installed. For users who are affected, he recommends that the following updates be uninstalled:

- Windows 10 1709: KB4493441
- Windows 10 1803: KB4493464
- Windows 10 1809: KB4493509
- Windows 10 1903: KB4495666

And... Bleeping Computer reports slow operation after Tuesday's patches
Paraphrasing from their report:

<quote> Users are reporting that after installing this week's Microsoft's April 2019 Patch Tuesday updates, Windows has suddenly become slow and programs are taking forever to open.

We have received emails and seen reports from users who have stated that this week's updates are also causing Windows to become very slow. The reports we have seen have been from users running Windows 7 and Windows 10. The issues that users are experiencing include Windows taking a long time to start or reboot, unable to start programs, lag in games, excessive disk activity, video streaming issues, and other similar problems. For example, in a comment at BleepingComputer a reader has stated that their Windows 10 computer has become extremely slow and that rebooting/starting Windows takes forever.

Users on Reddit [1, 2, 3, 4, 5, 6] and elsewhere [1,2] are also complaining that Windows has become very slow since installing the updates.

BleepingComputer's Lawrence Abrams writes: "Normally when a user has an antivirus program, Windows defender will disable its real time protection. It seems that for this user at least, Windows Defender is being enabled automatically even though the user had Avira installed on the machine.

Having two antivirus programs performing real-time protection could definitely cause slowdowns and other issues.

At this time there is nothing from Microsoft that states they are aware of the reported issues.

The only reference to Windows being slow since the updates is from a support article posted yesterday by Avira that is simply titled "Why does my system run very slow?". This article states that if Windows 10 has become slow, you should remove the KB4493509 update. For Windows 7 users, they state that you should remove KB4493472 and KB4493448 updates for Windows 7.

As these instructions are for users of their software, it may not apply to everyone.

BleepingComputer has contacted Microsoft to see if they are aware of these issues, but have not heard back as of yet.

Lawrence notes: It should be noted that I personally have no issues after installing these updates for Windows 10 and running ESET Nod32. If you are having issues, can you please leave a comment and let us know if Windows Defender is enabled along with your AV software or other things you have tried to resolve the issue.

Two hours after this was published the article was updated to add that ComputerWorld's Woody Leonhard also reported that he is seeing users have slowdown issues on Windows 10 after installing these updates.

And two hours after that, another update: BleepingComputer has been told by a source familiar with the matter, that these issues are being caused by conflicts between the recent updates and antivirus software. While Microsoft originally reported that the antivirus conflicts were only causing Windows to freeze, it appears that there could be other symptoms. You can read more about the antivirus conflicts at this article.

What good did April's patch Tuesday accomplish?
It repaired 74 security flaws including two actively exploited Windows 0-days!

That makes this the second month in a row that a pair of actively-exploited 0-days were patched.

The two zero-days patched are similar, both being elevation of privilege (EoP) vulnerabilities impacting the Win32K.DLL, a core component of the Windows OS Kernel.

They were discovered and responsibly reported by separate security teams: Alibaba Cloud Intelligence Security Team, and Kaspersky Lab. And Microsoft describes the two zero-days identically:

"An elevation of privilege vulnerability exists in Windows when the Win32k component fails to properly handle objects in memory. An attacker who successfully exploited this vulnerability could run arbitrary code in kernel mode. An attacker could then install programs; view, change, or delete data; or create new accounts with full user rights. To exploit this vulnerability, an attacker would first have to log on to the system. An attacker could then run a specially crafted application that could exploit the vulnerability and take control of an affected system. The update addresses this vulnerability by correcting how Win32k handles objects in memory."

No further details are available about the exploitation of the two vulnerabilities, except that they've been under active exploitation. However, Kaspersky has reported to Microsoft six Windows Win32k elevation of privilege zero-days in the past six months, all which we found being exploited by a nation state affiliated hacking group, so we might safely assume that the one Kaspersky found is number seven in the hit parade.

Aside from the Windows 0-days, not surprisingly, among the remaining 72 flaws fixed there were three Microsoft Office Access Connectivity bugs (CVE-2019-0824, CVE-2019-0825, CVE-2019-0827) that can allow attackers to execute code on vulnerable systems. All which can be exploited remotely. And another code execution bug (CVE-2019-0853) impacts the Windows GDI+ component when parsing EMF files. Since this vulnerability can be exploited merely by convincing users to visit a website or by emailing users malicious files, it's a serious flaw that needs patching.

I considered going into further detail, but the 74-entry list of every flaw, where it is and what it does, is mind numbing. Suffice to say that everything Microsoft publishes has numerous bugs of varying severity and that waiting long to apply last week's suite of updates (so long as you don't use 3rd-party A/V) would not be advisable.

And not to be left behind, Adobe also released 40 patches last Tuesday
This was a large patch security update resolving a host of critical and important bugs. Adobe provided patches for Adobe Bridge CC, Adobe Experience Manager Forms, InDesign, Adobe XD, Adobe Dreamweaver, Adobe Shockwave Player, Adobe Flash Player, and Adobe Acrobat and Reader. The vulnerabilities fixed include some which can lead to arbitrary code execution problem, sensitive information disclosure, and remote code execution in the context of the current user.

When I read that Adobe's Shockwave Player was suffering from a total of seven serious security flaws -- all critical memory corruption issues exploitable for the purpose of executing arbitrary code -- I thought... Shockwave? You've got to be kidding me. We're no longer allowed to use Windows 95, 98, NT, 2000, or XP ... But someone, somewhere, is still using Shockwave? That just doesn't seem right.

Well, actually, I spoke too soon... Since these were the LAST UPDATES that shockwave will ever receive: https://helpx.adobe.com/shockwave/shockwave-end-of-life-faq.html

<quote> Effective April 9, 2019, Adobe Shockwave will be discontinued and the Shockwave player for Windows will no longer be available for download. Companies with existing Enterprise licenses for Adobe Shockwave continue to receive support until the end of their current contracts.</quote>

So, that's that for Shockwave.

Abode's Flash player had an out-of-bounds read and a use-after-free flaw fixed, either of which could result in data leaks or the execution of arbitrary code.

Adobe Acrobat and Reader also received a substantial update last Tuesday with a total of 21 security issues resolved: 10 leading to information disclosure and 11 other bugs that could be exploited to execute arbitrary code.

Google considering automatically blocking "high risk" downloads
http://lists.w3.org/Archives/Public/public-webappsec/2019Apr/0004.html

Emily Stark's posting on the W3.ORG list was titled: "Blocking high-risk non-secure downloads"

Emily wrote… “Hi webappsec friends,

Over in Chrome land, we've been considering how to drive down non-secure downloads, particularly high-risk ones like executables. I wanted to see if other browsers would be interested in joining us on this adventure.

We want to achieve the right balance between compatibility/user-disruption and security improvements, so we will likely start by treating certain high-risk downloads initiated from secure contexts as active mixed content and block them. We're still finalizing our metrics before we can share them publicly, but right now it's looking like it will be feasible to block a set of high-risk filetypes (executables and archives as determined by the Content-Type header or sniffed mime-type). We will likely focus on protecting desktop users because Android and Safe Browsing already provide protection against malicious APKs.

https://source.android.com/security/reports/Google_Android_Security_2018_Report_Final.pdf

We're not planning to focus on non-secure downloads initiated from non-secure contexts at the moment, because users at least see the "Not Secure" omnibox badge on those pages.

Feedback welcome! / Thanks, Emily

In a follow-up reply to someone's query about which types Google was considering, Emily wrote: We are looking at exes, dmgs, and crxs as executables, and zip/gzip/rar/tar/bzip/etc. as archives.

In response to a query from ZDNet, a Mozilla spokesperson said: "We are interested in exploring these ideas further in conversation with Google and other interested parties. The general idea aligns with the steps we have previously taken to protect users from insecurely delivered content."

The idea of blocking non-secured queries from secured pages is not new. Our browsers already have many provisions for special-casing a drop in security. As with anything, such a move would tend to break some sites, but unlikely anything mainstream, and those sites can probably arrange to use secure downloads.

I would chalk this up to another incremental step toward the deprecation of non-HTTPS connections on the web.

Russia's Roskomnadzor finally lowered the boom on Facebook!!

We've covered this pending and growing issue previously.

As we noted at the time, last December, Russian Internet watchdog Roskomnadzor sent notifications to both Twitter and Facebook asking them to provide information about the location of servers that store the personal data of its citizens. Remember that Roskomnadzor – also known as the Federal Service for Supervision in the Sphere of Telecom, Information Technologies, and Mass Communications – is the Russian telecommunications watchdog that runs a huge blacklist of websites banned in Russia.

Though the social media platforms were given one month to reply, they choose too stick to their guns and to NOT disclose this information. And as a result, Moscow's Tagansky District Court imposed a whopping 3,000 rubles fine on Twitter last week and the same on Facebook today.

And what is 3,000 rubles in US dollars? Well... Brace yourself... It's $47. Yup.

That fine was the minimum that Russian courts can impose on companies for violating Article 19.7 of the Administrative Code of the Russian Federation, i.e., failure to provide information. The maximum amount of the fine under this article is ... Wait for it ... 5,000 rubles or $78.

Twitter and Facebook are not off the hook, however, since Russia law does give them the ability to completely ban non-complying social media companies as they successfully banned LinkedIn back at the end of 2016.

“Taj Mahal”

https://securelist.com/project-tajmahal/90240/

Kaspersky Lab named the massive APT framework suite "TajMahal" because the stolen data was transferred to the attackers' C&C server in an XML file named TajMahal.

Kaspersky describes this as a state-of-the-art, high-tech, modular-based malware toolkit that not only supports a vast number of malicious plugins for distinct espionage operations, but also comprises never-before-seen and obscure tricks. Evidence shows that the system has been in active operation for at least 5 years, but that it managed to remain undetected until just recently. Malware samples they examined suggest the cyberespionage group behind the attack has been active since at least August 2014.

The system pinged Kaspersky's radar late last year the attackers used it to spy on the computers of a diplomatic organization belonging to a Central Asian country whose nationality and location have not been disclosed.

Kaspersky wrote:

‘TajMahal’ is a previously unknown and technically sophisticated APT framework discovered by Kaspersky Lab in the autumn of 2018. This full-blown spying framework consists of two packages named ‘Tokyo’ and ‘Yokohama’. It includes backdoors, loaders, orchestrators, C2 communicators, audio recorders, keyloggers, screen and webcam grabbers, documents and cryptography key stealers, and even its own file indexer for the victim’s machine. We discovered up to 80 malicious modules stored in its encrypted Virtual File System, one of the highest numbers of plugins we’ve ever seen for an APT toolset.

Just to highlight its capabilities, TajMahal is able to steal data from a CD burnt by a victim as well as from the printer queue. It can also request to steal a particular file from a previously seen USB stick; next time the USB is connected to the computer, the file will be stolen.

TajMahal has been developed and used for at least the past five years. The first known ‘legit’ sample timestamp is from August 2013, and the last one is from April 2018. The first confirmed date when TajMahal samples were seen on a victim’s machine is August 2014.

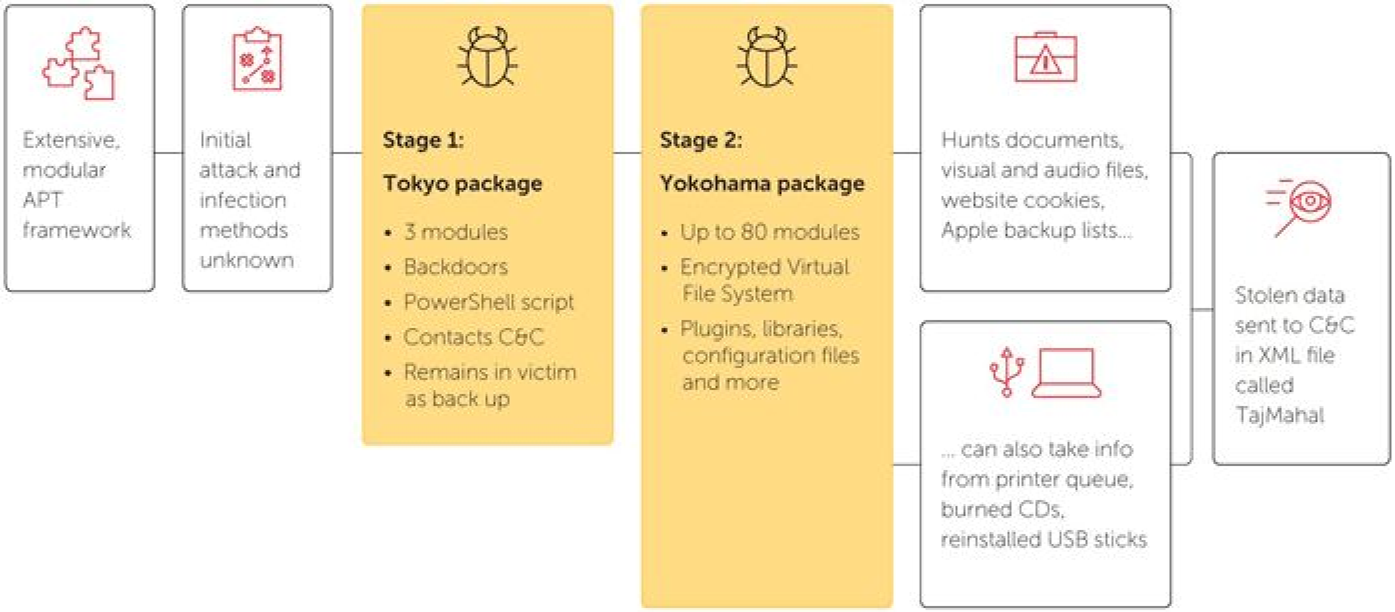
Extensive, modular APT framework
Initial attack and infection methods unknown
Stage 1:
Tokyo package
• 3 modules
• Backdoors
• PowerShell script
• Contacts C&C
• Remains in victim as back up
Stage 2:
Yokohama package
• Up to 80 modules
• Encrypted Virtual File System
• Plugins, libraries, configuration files and more
Hunts documents, visual and audio files, website cookies, Apple backup lists...
... can also take info from printer queue, burned CDs, reinstalled USB sticks
Stolen data sent to C&C in XML file called TajMahal

Kaspersky's report concluded:

The TajMahal framework is an intriguing discovery that’s of great interest, not least for its high level of technical sophistication, which is beyond any doubt. The huge amount of plugins that implement a number of features is something we have never before seen in any other APT activity. For example, it has its own indexer, emergency C2s, is capable of stealing specific files from external drives when they become available again, etc.

The question is, why go to all that trouble for just one victim? A likely hypothesis is that there are other victims we haven’t found yet. This theory is reinforced by the fact that we couldn’t see how one of the files in the VFS was used by the malware, opening the door to the possibility of additional versions of the malware that have yet to be detected.

Miscellany

Firefox: "Auto Tab Discard" - removes memory occupied by non-current tabs.

SpinRite

Andy Weaver in Bath, UK
Subject: Spinrite
Date: 10 Apr 2019 23:47:28
:
Really enjoy Security Now - the subject coverage is always interesting and just the right level of technical detail for my level of tech knowledge.

Wondering about the long-awaited Spinrite update. If I buy now, will I be entitled to the promised updated features when released?

Many thanks. Live long and prosper!
Andy / Bath / UK

Closing The Loop

Roy in Israel
Subject: Another reason to use a password manager or SQRL
Date: 16 Apr 2019 00:48:43
:
Hi Steve,

I have been listening to the show for a few months now.
I found this out when one of our support engineers opened a bug about the ability to retrieve the password in our service login page using a simple inspect and replace type method within chrome.
Of course, there was nothing we can do about it but I wanted to share this anyway so people understand how dangerous is the simple password manager which is embedded within our browser.

See here:
https://www.maketecheasier.com/see-password-in-browser/

Thanks for a great show,
Roy

Scott in Boston
Subject: Security Now feedback: uBlockOrigin disables <a> ping attribute by default
Date: 15 Apr 2019 16:43:06
:
Hi Steve,
Per your story last week on the formalization of the ping attribute in the HTML spec, it looks like uBlock origin already blocks sending that info.

It's in the Settings tab under Privacy as "Disable hyperlink auditing"
"Checking this will prevent hyperlink auditing. Hyperlink auditing is best summarized as "phone home" feature (or more accurately "phone anywhere") meant to inform one or more servers of which links you click on (and when)."
The explanatory link goes to a page that defines "hyperlink auditing" as encompassing the ping attribute as well as a DOM method called navigator.beacon, so I believe both of those are blocked by uBlockOrigin when Disable hyperlink auditing is checked

Richard in York, UK
Subject: Copy as path
Date: 10 Apr 2019 23:25:46
:
Steve

I've been listening to SN for a few years and absolutely love it. More than that, it's been super useful as well. But in 709, you have revolutionised my working on Windows. I'm forever needing to type in/copy file paths for various things and I can hardly believe that 'copy as path' has always been there and I didn't know about it! Thank you so much for letting the world know about that! It seems like a very little thing but it is just so incredibly quick and useful. *mind blown*

Anonymous Sender Somewhere
Subject: SN 709
Date: 14 Apr 2019 00:24:21
:
Steve,

Thank you very much for the Windows Explorer tip in SN 709. I find it very useful.

I have also observed that when use Shift+"Right Click", the "Send To" menu is expanded, which is an additional bonus.

Regards.

DragonBlood

WPA3 is beginning to come under scrutiny and implementation faults are arising.

Mathy Vanhoef, a researcher who was at the University of Leuven (KU Leuven) two years ago discovered and revealed a severe flaw in the Wi-Fi Protected Access II (WPA2) protocol that we're all still using today. He named the attack "KRACK" for Key Reinstallation Attack.

Today, Mathy is at New York University Abu Dhabi and working with another researcher at Tel Aviv University and KU Leuven. Their new research paper is titled: Dragonblood: A Security Analysis of WPA3's SAE Handshake".

As we've mentioned here in our initial preliminary discussion of WPA3, the WPA3-Personal protocol replaces WPA2's Pre-shared Key (PSK) with a protocol called "Simultaneous Authentication of Equals" or (SAE). It is intended to provide more robust password-based authentication. This SAE protocol is also known as "Dragonfly" and it appears to contain a number of fundamental design flaws which expose users to password partitioning attacks.

So... DragonBlood are the newly discovered attacks on the DragonFly protocol.

Abstract

The WPA3 certification aims to secure Wi-Fi networks, and provides several advantages over its predecessor WPA2, such as protection against offline dictionary attacks and forward secrecy. Unfortunately, we show that WPA3 is affected by several design flaws, and analyze these flaws both theoretically and practically. Most prominently, we show that WPA3's Simultaneous Authentication of Equals (SAE) handshake, commonly known as Dragonfly, is affected by password partitioning attacks. These attacks resemble dictionary attacks and allow an adversary to recover the password by abusing timing or cache-based side-channel leaks. Our side-channel attacks target the protocol's password encoding method. For instance, our cache-based attack exploits SAE's hash-to-curve algorithm. The resulting attacks are efficient and low cost: brute-forcing all 8-character lowercase passwords requires less than $125 in Amazon EC2 instances. In light of ongoing standardization efforts on hash-to-curve, Password-Authenticated Key Exchanges (PAKEs), and Dragonfly as a TLS handshake, our findings are also of more general interest. Finally, we discuss how to mitigate our attacks in a backwards-compatible manner, and explain how minor changes to the protocol could have prevented most of our attacks.

Introduction

TheWi-Fi Alliance recently announced WPA3 as the more secure successor of WPA2. Unfortunately, it was created without public review, meaning experts could not critique any of WPA3's new features before they were released. Moreover, although the new handshake of WPA3 was designed in an open manner, its security guarantees are unclear. On one hand there is a security proof of a close variant of WPA3's handshake, but on the other hand another close variant of the handshake received significant criticism during its standardization. These issues raise the question whether WPA3 is secure in practice. We remark that WPA3 does not define new protocols, but in-stead mandates which existing protocols a device must support.This means WPA3 is not a specification, but a certification. Put differently, devices can now become WPA3-certified, which assuresthey implement certain protocols in an interoperable manner. The only novelty in the WPA3 certification is a transition mode where WPA2 and WPA3 are simultaneously supported for backward compatibility. Although WPA3 follows recommended practice by using existing standards, we believe more openness to alternative protocols could have increased its security.

In this paper we perform a security analysis of WPA3's Simul-taneous Authentication of Equals (SAE) handshake. This hand-shake is designed to prevent dictionary attacks, and constitutes thebiggest improvement over WPA2. We systematically analyzed itssecurity by reading specifications, inspecting formal proofs, and au-diting open-source implementations. This analysis revealed several design and implementation flaws. For instance, when verifying the assumptions made by the formal proof of the SAE handshake, we discovered both timing and cache-based side-channel vulnerabilities in its password encoding method. We empirically confirmed all our findings against both open source and recently-released proprietary implementations of WPA3.

All combined, our work resulted in the following contributions:

- We provide a self-contained and high-level description of WPA3 and its SAE handshake.
- We show that the anti-clogging mechanisms of SAE is unable to prevent denial-of-service attacks. In particular, by abusing the overhead of SAE's defenses against already-known side-channels, a resource-constrained device can overload the CPU of a professional Access Point (AP).
- We present a dictionary attack against WPA3 when it is operating in transition mode. This is accomplished by trying to downgrade clients to WPA2. Although WPA2's 4-way handshake detects the downgrade and aborts, the frames sent during the partial 4-way handshake provide enough information for a dictionary attack. We also present a downgrade attack against SAE, and discuss implementation-specific downgrade attacks when a client improperly auto-connects to a previously used WPA3-only network.
- We empirically investigate the feasibility of timing attacks against WPA3's SAE handshake. This confirms timing attacks are possible and leak info about the password.
- We present a novel micro-architectural cache-based side-channel attack against the SAE handshake. This attack leaks information about the password being used. Our attack even works against hash-to-curve algorithm implementations that include countermeasures against side-channel attacks. This type of attack against hash-to-curve algorithms is of independent interest due to current standardization efforts surrounding hash-to-curve methods.
- We show both theoretically and empirically how the recovered timing and cache info can be used to perform an offline password partitioning attack. This enables an adversary to recover the password used by the victim.

10 CONCLUSION AND RECOMMENDATIONS

In light of our presented attacks, we believe that WPA3 does not meet the standards of a modern security protocol. Moreover, we believe that our attacks could have been avoided if the Wi-Fi Alliance created the WPA3 certification in a more open manner. Notable is also that nearly all of our attacks are against SAE's password encoding method, i.e., against its hash-to-group and hash-to-curve algorithm. Interestingly, a simple change to this algorithm would have prevented most of our attacks. In particular, the peer's MAC addresses can be excluded from SAE's password encoding algorithm, and instead included later on in the handshake itself. This allows the password element to be computed offline, meaning an adversary can no longer actively trigger executions of the password encoding method. Moreover, this would mean that for a given password, the execution time of the password encoding method would always be identical, limiting the amount of information being leaked. Surprisingly, when the CFRG was reviewing a minor variant of Dragonfly, they actually discussed these type of modifications [48,49,69,80]. However, to our surprise, this change was not incorporated into any of the Dragonfly variants. We also conjecture that resource-constrained devices may not implement all the side-channel countermeasures, as these maybe too costly on lightweight processors. Additionally, correctly implementing our suggested backwards-compatible side-channel countermeasures is non-trivial. This is worrisome, because security protocols are normally designed to reduce the chance of implementation vulnerabilities. Finally, we believe that a more open process would have prevented (or clarified) the possibility of downgrade attacks against WPA3-Transition mode. Nevertheless, although WPA3 has its flaws, we still consider it an improvement over WPA2.

It's also important for us to remember that 14 years ago, when WPA2 was happening, driven by the total catastrophe that was WPA, MOST of the traffic being carried over WiFi was UN-encrypted. Those were the quaint and heady days of Firesheep and promiscuous traffic sniffing where web pages only briefly used HTTPS while their user was logging in to protect their password, then switch back to HTTP, thus exposing their browser's session cookie... making impersonation trivial. And back then eMail was almost never encrypted.

Today, where HTTPS and TLS have become virtually ubiquitous, the role of WiFi encryption has shifted a bit. It is no longer needed as much to protect our web and eMail traffic, but it remains crucial for keeping wireless bad guys out of our networks where we increasingly have IoT and an ever increasing number of "servers" of various sorts. If I were responsible for any sort of enterprise deployment, all of our WiFi access points would be on their own network segment with NO access to the corporation's vital innards. The risk is just too great.

From the Wi-Fi Alliance: Security Update April 2019
https://www.wi-fi.org/security-update-april-2019

As with any technology, the robust security research necessary to remain ahead of emerging threats will occasionally uncover new vulnerabilities. Security researchers identified vulnerabilities in a limited number of early implementations of WPA3™-Personal and immediately brought their discovery to the Wi-Fi® industry. There is no evidence of the vulnerability being used against Wi-Fi users maliciously, and Wi-Fi Alliance® has taken immediate steps to ensure users can count on WPA3-Personal to deliver even stronger security protections.

- Wi-Fi CERTIFIED WPA3-Personal now includes additional testing within our global certification lab network to encourage greater adoption of recommended practices

- Wi-Fi Alliance is broadly communicating details on these vulnerabilities and implementation guidance to device vendors as the industry begins to bring WPA3-Personal to market

These issues can be resolved through a straightforward software update – a process much like the software updates Wi-Fi users regularly perform on their mobile devices. WPA3-Personal is in the early stages of deployment, and the small number of device manufacturers that are affected have already started deploying patches to resolve the issue. The software updates do not require any changes that affect interoperability between Wi-Fi devices. Users can refer to their device vendors' websites for more information.

As always, Wi-Fi users should ensure they have installed the latest recommended updates from device manufacturers. Security is and always will be a dynamic endeavor, and Wi-Fi Alliance will continue to maintain strong security protections for Wi-Fi users through its Wi-Fi CERTIFIED™ program.

~30~