Security Now! #712 - 04-30-19
Credential Stuffing Attacks

This week on Security Now!

This week we look at more privacy fallout from our recent coverage of Facebook and Google, we examine the uptake rate of recent Windows 10 feature releases, we finally know the source of the A/V troubles with the April patch Tuesday updates, we look at the NIST's formal fuzzing development, consider the source of a massive and ongoing database data leak involving more than half of all American households, we note that Windows Insiders are already finding that their systems won't update to the May 2019 feature update, and we address the concerns of United Airlines passengers who have noticed and been understandably upset by seatback cameras pointing at them. Finally, we have the "Cranky Old Guy Tip of the Week", we touch on a bit of miscellany, then take a look at what many in the security industry are watching with concern: The large and emerging threat of website credential stuffing attacks.

Current Windows 10 Release Distribution

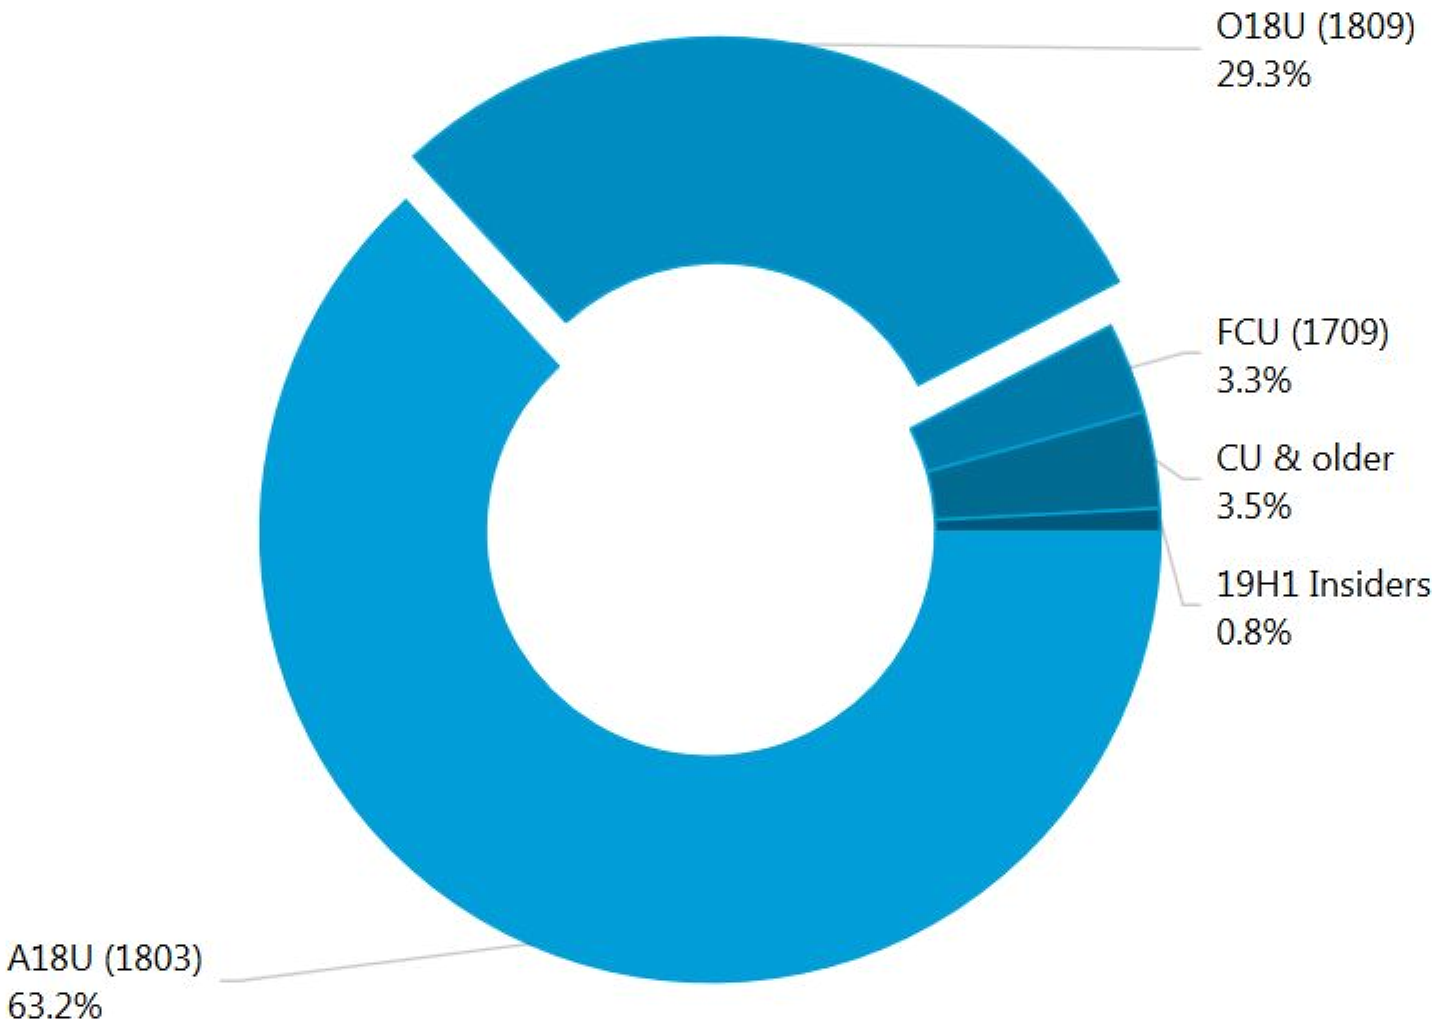

Security News

Facebook was fined $47 dollars by Russia, but may be facing a $5 Billion fine related to the Cambridge Analytica data disclosure.

As we noted last week, Russia has fined Face 3,000 rubles which is $46.53. But it appears that Facebook may be expecting a fine of up to $5 billion from the US Federal Trade Commission as a consequence of their year-long investigation of the 50 million Facebook user records that were turned over to Cambridge Analytica without their users' knowledge or consent.

Although no mention was made of this during Facebook's earnings call last Wednesday, in the release of its first quarter earnings, Facebook said that it would be "setting aside" $3 to $5 billion as a contingency expense "in connection with" the FTC's investigation of its user-data practices.

The written release states: "In the first quarter of 2019, we reasonably estimated a probable loss and recorded an accrual of $3 billion in connection with the inquiry of the FTC into our platform and user-data practices, which accrual is included in accrued expenses and other current liabilities on our condensed consolidated balance sheet. We estimate that the range of loss in this matter is $3 billion to $5 billion. The matter remains unresolved, and there can be no assurance as to the timing or the terms of any final outcome." FWIW, this sort of statement is commonplace (even though the amount may not be) in these sorts of earnings reports to shareholders. To avoid shareholder lawsuits, which are all too common, painting a somewhat bleak downside is an effective CYA measure.

Last March, 2018, the FTC announced that it was launching an investigation into Facebook's data-privacy practices, triggered by the Cambridge Analytica revelations. The FTC is specifically investigating whether Facebook has violated a consent decree from 2011 that requires the social network to receive explicit permission from users before sharing their data with third parties.

In a conference call last Wednesday afternoon with analysts, Mark Zuckerberg said that he plans to forge ahead with what he calls a privacy-focused social network. He did not refer to the FTC fine or the company's recent data-handling problems. However, the company's CFO, David Wehner, commented that "we anticipate ad targeting-related headwinds will be more pronounced in the second half of 2019" — you've GOT to be kidding me! "ad targeting-related headwinds will be more pronounced" ??? Sheesh. Talk about corporate doublespeak. Wehner also noted that he expects new regulations to affect the company's business going forward.

Uh huh.

To give us some sense of perspective, on first-quarter 2019 revenue of $15.08 billion, Facebook reported a profit of $2.43 billion, which is 85 cents per share. This is a drop from last year's same-quarter per share revenue of $1.69 net income on $11.97 billion in revenue for the year-ago quarter. This drop is due, in part, to the charge for the expected fine. Despite all the recent bad news, Facebook reported that both monthly and daily active users were up 8 percent from last year, coming in at 2.38 billion and 1.56 billion, respectively. And even in the immediate aftermath of the massive expected fine, Facebook's stock was up in after-hours trading.

SensorVault: Congress Asks Google 10 Questions On Its Location Tracking Database
https://energycommerce.house.gov/sites/democrats.energycommerce.house.gov/files/documents/Google.2019.4.23.%20Letter%20to%20Google%20re%20Sensorvault.CPC_.pdf

From: *the House of Representatives Committee on Energy and Commerce.*
To: *Mr. Sundar Pichai, CEO, Google.*

Dear Mr. Pichai:
We are writing in response to concerning reports about a massive database of Precise location information on hundreds of millions of consumers known inside Google as "SensorVault." According to recent reports, Google tracks and stores precise location information on a huge volume of consumers, including practically every consumer with an Android mobile device, in some cases storing information dating back to 2009.

The potential ramifications for consumer privacy are far reaching and concerning when examining the purposes for the SansorVault database and how precise location information could be shared with third parties.

First, according to the reports, Google collects precise location information in numerous ways including from the location history function on Android Phones, Google searches, and Google apps that have location enabled. Second, precise location information is reportedly collected even when people are not making calls or using apps, which enabled Google to track the "whole pattern of life" of an individual. Finally, Google reportedly never destroys any of the precise location information it captures in the SensorVault database, and has therefore compiled an extraordinarily details picture of the movements and whereabouts of a vast number of consumers stretching back more than a decade.

As part of the Committee's ongoing commitment to protect the privacy of the American people and to understand the benefits and risks of various data collection and use practices, we would like to know the purposes for which Google maintains the SensorVault database and the extent to which Google shares precise location information from this database with third parties. To that end, please provide answers to the following questions by May 7, 2019:

1. What information does Google store in the SensorVault database and for what purposes does Google use this information? Please described each use in detail. If the types of information in the database or the purposes for which such information is used have changed over time, explain any such changes in chronological order.

2. Please describe which affiliates/subsidiaries of Alphabet have access to or use the data or analytics derived from the data in the SensorVault database.

3. Does Google maintain other databases of precise location information? If so, how does the SensorVault database differ from other such databases and how is the data from such database used?

4. Who I able to access the information in the SensorVault database? Include in your response both the number of Google employees with access to the SensorVault database and the roles and responsibilities of any persons with access.

5. What are the sources from which Google collects the information maintained in the SensorVault database and any other database identified in response to question 3? Specifically, describe any Google services, mobile applications, devices, or any other means through which Google obtains the information.

   a. If consumers are required to "opt in" to the collection of precise location information, describe with specificity how consumers "opt in" to the service, and the notice provided to such consumers about the purposes for which Google collects the information. If any such “opt in” has changed over time, describe those changes.

   b. If consumers may "opt out" of the collection of precise location information, describe with specificity how consumers "opt out" and the notice provided to such consumers about the purposes for which Google collects the information. If any "opt out" has changed over time, describe those changes.

6. To the extent that a consumer has requested that precise location data not be shared with Google, through opt-outs or other mechanisms, do Android phone continue to collect precise location data on the device or store precise location information on the device? If so, to the extent that the consumer subsequently allows location data to be shared with Google, is that formerly stored precise location information transmitted to Google? Under what other circumstances would the device continue to transmit precise location information to Google when a consumer has requested that precise location information not be shared with Google?

7. How accurate is the precise location information stored in the SensorVault database? Include in your response both the accuracy of the precise location information and the accuracy of attributing such information to a single individual.

8. What controls, if any, does Google provide to consumers to limit or revoke Google's access to the information stored in the SensorVault database? Does Google provide consumers a means to delete data stored in SensorVault? If so, describe in detail how consumers may do so.

9. What is Google's retention policy with respect to precise location information stored in the SensorVault database? For what purpose(s) does Google maintain precise location information going back to 2009.

10. Does Google share, sell, license, or otherwise disclose precise location information (including deidentified data) from the SensorVault database with any third parties other than law enforcement? If so, identify the types of businesses receiving such information and the purpose for disclosing any such information. If the data is de-identified, provide a description of how it is de-identified.

In addition to providing the written responses requested, please make arrangements to provide Committee staff with a briefing on these topics to occur no later than May 10, 2019. Thanks you for your attention. If you have any questions, please contact Lisa Goldman of the Majority Staff at (202) 225-2927 and Melissa Froelich of the Minority Staff at (202) 226-3641.

Sincerely, Frank Pallone, Jr., Chairman / Greg Walden, Ranking Member
Jan Schakowsky, Chair, Subcommittee on Consumer Protection and Commerce
Cathy McMorris Rodgers, Ranking Member, Subcommittee on Consumer Protection and Commerce

Windows 10 shows a slow October 2018 uptake...
https://reports.adduplex.com/#/r/2019-04

A Lithuania company called "AdDuplex" calls itself the largest cross-promotion network for Windows Store apps and games empowering developers and publishers to promote their apps for free by helping each other. It was established eight years ago in 2011 and today more than 10,000 apps actively use AdDuplex to gain increased visibility and monetize their ad space.

Just last Friday, April 26th, the version of Windows being run on approximately 5,000 Windows store apps incorporating the v.2 and higher of the AdDuplex SDK was collected. This was a sampling of more than 100,000 Windows 10 machines.

63.2% of users are running Windows 10 April 2018 Update (version 1803).
29.3% are using Windows 10 October 2018 Update (version 1809).
3.5% are using the Creators Update or older.
3.3% are using the Fall Creators Update (version 1709).
0.8% are using 19H1 Windows Insider builds (version 1903).

What's really interesting is this chart of the comparative historical uptake rate of all previous Win10 releases:

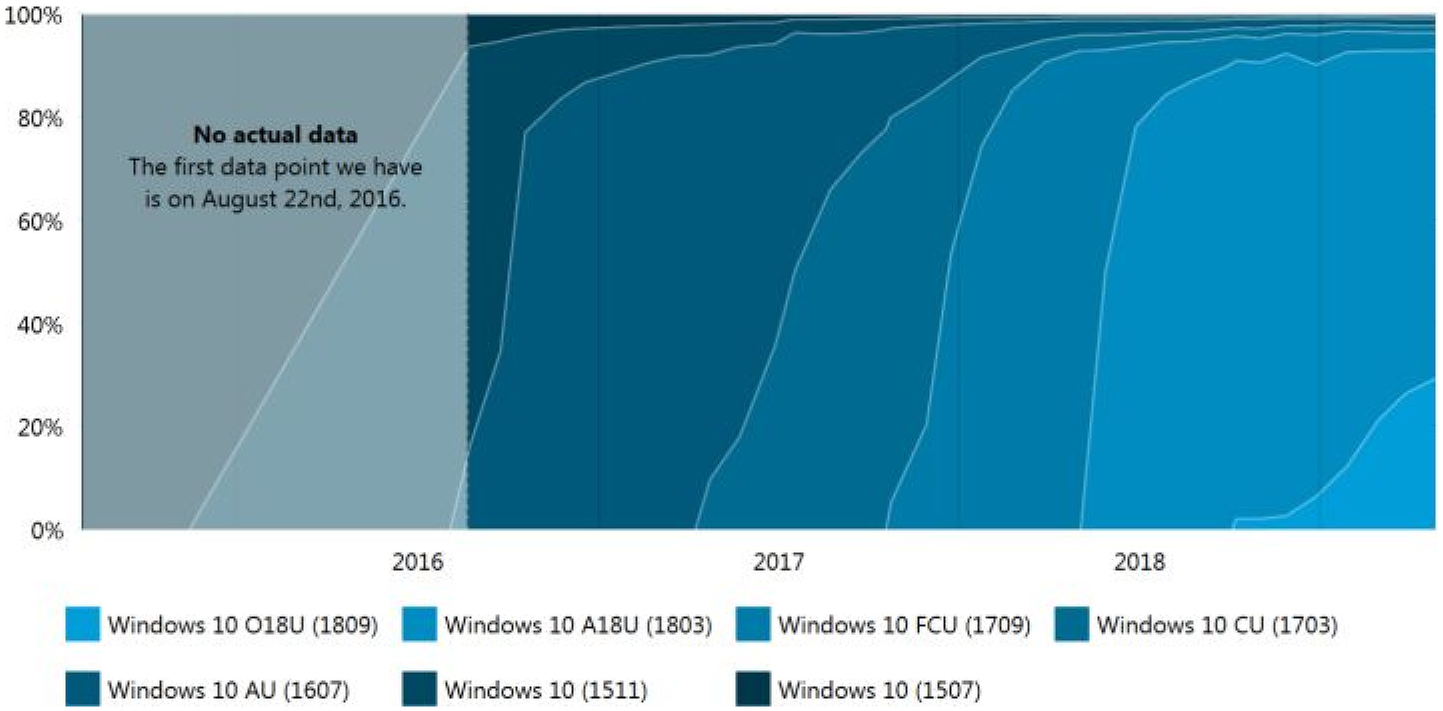

The leading edge slope of each of those curves shows the uptake rate at from the release point of each version. Certainly, last October 2018's Feature Update was a catastrophe. But it's now been officially released and made available. All of my Win10 machines which are NOT running the LTSC version of Win10 have self-updated to 1809. So how do we account for this. It has been suggested that perhaps users are gun shy and are wishing to hold off. But I'm 100% certain that the typical Windows 10 user has no idea about what's been going on behind the scenes with the past six months of the Win10 update debacle.

So something must be throttling the move to 1809. I suspect there will be some users who have used the advanced windows update settings to disable the feature update. And it might be that Microsoft released 1809 mostly to save face by releasing it at all. Because the next big feature release, the May 2019 Update - version 1903 - is due shortly. So maybe Microsoft plans to mostly just leapfrog over 1809 and ever expose the remaining 2/3rds of its Win10 user to it at all. And the last explanation is that many Win10 machines are inside corporate enterprise IT control and they are certainly aware of what a disaster 1809 was. So corporate users may well be hanging back and stacking with 1803 which they know at least works.

It's going to be very interesting to watch the adoption rate of 1903 once it hits shortly. Since it appears that AdDuplex publishes monthly reports, we'll be able to check back and watch the fun!

And speaking of trouble being caused by Windows update...

Bleeping Computer's Lawrence Abrams did some creative sleuthing and appears to be the first person to figure out and clearly disclose exactly which of the April 2019 updates was responsible for so many of the enterprise A/V systems having trouble earlier this month.

As we know, in eight separate support articles for the April Windows updates, Microsoft explained that if users have enterprise antivirus software from Avast, Avira, McAfee, or Sophos installed, Windows may "become unresponsive upon restart after installing."

Lawrence had observed that both McAfee and Avast had made passing references to CSRSS being related to the problems.
CSRSS is the Client Server Run-Time Subsystem within Windows. So he looked through the list of updates released on April 9th and noticed that a security update was released for "CVE-2019-0735 | Windows CSRSS Elevation of Privilege Vulnerability".

This security update stated: "An elevation of privilege vulnerability exists when the Windows Client Server Run-Time Subsystem (CSRSS) fails to properly handle objects in memory. An attacker who successfully exploited this vulnerability could run arbitrary code. 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 to take control of an affected system. The update addresses the vulnerability by correcting how the Windows CSRSS handles objects in memory."

That appears to have been the culprit that caused those 3rd-party A/V tools so much trouble after the April 9th updates. The vulnerability was privately disclosed to Microsoft and was not known to be actively exploited in the wild.

NIST has gone Fuzzing
"NIST" is the US National Institute of Standards and Technology.

After more than 20 years of steady improvement, the NIST believes that it has reached an important milestone in the development is a technology it calls Combinatorial Coverage Measurement (CCM). It appears to be what hackers commonly refer to as "Fuzzing."

A a part of a larger research toolkit called Automated Combinatorial Testing for Software (ACTS), CCM is an algorithmic approach used to test software for interactions between input variables that might cause unexpected failures.

This is potentially useful for testing and hardening complex and important mission-critical systems such as aircraft, cars and power plants where problems can be life-threatening. Many complex systems operating in real must continually receive inputs from large arrays of sensors which might generate unexpected conflicts the software can’t resolve. The systems' designers work to identity and counteract these problems by modelling as many interactions as they can before the software is used in the real world, but what if something that COULD happen is missed? This is where ACTS and CCM come in.

However, we know what happens when systems grow. In the same way that adding each binary bit to a number doubles the number of possible states in a system, adding another sensor input explodes the possible states for a system exponentially. So the problem becomes on of modelling enough interactions from enough variables to spot all the possible combinations that might lead to an issue.

Due to 20 years of work, the solutions have been improving. But as the result of a collaboration with the University of Texas, Austria’s SBA Research (which describes itself as the research center for information security in Austria), and Adobe (one of several large companies using the toolkit), NIST thinks that the 2019 revision of CCM has made a significant leap forward. With the help of a new algorithm developed by SBA Research, NIST’s tool has jumped from being able to model a few hundred variables to up to 2,000 from five-way combinations of inputs. This means that it can now be practically applied to much more complex systems than previously.

https://www.nist.gov/news-events/news/2019/04/nist-tool-enables-more-comprehensive-tests-high-risk-software

The NIST's announcement last Tuesday was headlined: " NIST Tool Enables More Comprehensive Tests on High-Risk Software" with the subhead: Updated “combinatorial testing” tool could reduce potential errors in safety-critical applications.

NIST mathematician Raghu Kacker said that CCM represents a substantial improvement to the ACTS toolkit since its last major addition in 2015. “Before we revised CCM, it was difficult to test software that handled thousands of variables thoroughly,” Kacker said. “That limitation is a problem for complex modern software of the sort that is used in passenger airliners and nuclear power plants, because it’s not just highly configurable, it’s also life critical. People’s lives and health are depending on it.”

So I suppose we ought to call this "high-end formal fuzzing" HEFF?

vpnMentor reports: Unknown Data Breach Exposes 80 Million US Households
https://www.vpnmentor.com/blog/report-millions-homes-exposed/

Hosted by a Microsoft cloud server, an exposed 24 GB database includes the number of people living in each household with their full names, their marital status, income bracket, age, and more.

The database seems to itemize households rather than individuals. It includes:

- Full addresses, including street addresses, cities, counties, states, and zip codes
- Exact longitude and latitude
- Full names, including first, last, and middle initial
- Age
- Date of birth
- Title
- Gender
- Marital status
- Income
- Homeowner status
- Dwelling type

In scanning through interesting news of the past week for this podcast I am continually seeing articles about data breaches. We all know that they are all happening all the time. It just gets to be another form of Internet background radiation. For example, three recent headlines that I skipped over were:

- Medical Information of Almost 150K Rehab Patients Exposed
- 540 Million Facebook Records Leaked by Public Amazon S3 Buckets
- 257K Legal Documents Leaked By Unprotected Elasticsearch Server

Normally it's possible to identify the owner of the data and then notify them that they've left their fly open.

But in this case, the research team at vpnMentor has been unable to figure out who all this data might belong to. And what's more, weighing in at 80 million households means that it's over half of the households in the U.S.

Here's what the vpnMentor guys explain, and their call for help to the public -- that's us:

The research team is currently undertaking a huge web mapping project. They use port scanning to examine known IP blocks. This reveals open holes in web systems, which they then examine for weaknesses and data leaks.

Usually, researchers suspect where the leak is coming from. They can then examine the database to confirm its identity.

We then reach out to the database’s owner to report the leak, and where possible, alert the people affected. This helps build a safer and more protected internet.

Although we investigated the database online, we didn't download it. Our researchers felt that downloading it would be an ethical breach, as they would then illegally own personally identifiable data sets without peoples' consent.

This time, it's different. The database that the team discovered includes identifying information for more than 80 million households across the United States. As most households include more than one resident, the database could directly impact hundreds of millions of individuals.

vpnMentor is calling on the public to help identify the database and close the leak.

Unlike previous leaks we've discovered, this time, we have no idea who this database belongs to. It's hosted on a cloud server, which means the IP address associated with it is not necessarily connected to its owner.

The data includes uniform entries for more than 80 million households, making it almost impossible to narrow down. The only clue we found lay in people's ages: despite searching thousands of entries, we could not find anyone listed under the age of 40.

Interestingly, a value for people's income is given (however, we don't know if it's a code for an internal ranking system, a tax bracket, or an actual amount).

This made us suspect that the database is owned by an insurance, healthcare, or mortgage company. However, information one may expect to find in a database owned by brokers or banks is missing. For example, there are no policy or account numbers, social security numbers, or payment types.
Help Us Identify this Database

We want to contact this database's owners and let them know that their data logs are exposing millions of households.

Help us solve the riddle:

What service is used by 80 million homes across the US – but only the US – and only by people over 40? What service would collect your homeowner status and dwelling type but not your social security number? And what service records that you're married but not how many children you have?

If you can help us identify this database or know who owns it, please contact us at info@vpnmentor.com. The 80 million families listed here deserve privacy, and we need your help to protect it.

And speaking of the May 2019 Windows 10 feature update...
Microsoft is explaining to those on the Windows Insider Program who are trying to apply the May 2019 feature update that: "*This PC can't be upgraded to Windows 10 because it has a USB device or SD card attached.*"

Yes, you heard that correctly. I'm not kidding. Microsoft's exact text reads:

> If you are part of the Window Insider Program, and you are trying to upgrade to the May 2019 Update for Windows 10 (Windows 10, version 1903), you may experience an upgrade block and receive the following message:
>
> This PC can't be upgraded to Windows 10.
>
> An external USB device or SD memory card that is attached to the computer could cause inappropriate drive reassignment on Windows 10-based computers during the installation of the May 2019 update. For this reason, these computers are currently blocked from receiving the May 2019 Update. This generates the error message that is mentioned in the "Symptoms" section if the upgrade is tried again on an affected computer.
>
> Example: An upgrade to the May 2019 Update is tried on a computer that has a thumb drive inserted into a USB port. Before the upgrade, the device would have been mounted in the system as drive G based on the existing drive configuration. However, after the upgrade, the device is reassigned a different drive letter. For example, the drive is reassigned as drive H.
>
> Note: The drive reassignment is not limited to removable drives. Internal hard drives can also be affected.

https://support.microsoft.com/en-us/help/4500988/windows-update-blocked-for-windows-10-insider-program

This will presumably be fixed before the big roll-out of next month's release, otherwise the adoption of this release may not proceed any faster than the last one. :(

United Airlines purchased seat back displays with cameras

Passengers became quite upset at the idea of a camera staring at them throughout their flight. This is exactly where you want to have a physical mechanical slide door to shut the camera labeled with clear instructions, and the inter-flight cleaning crew should be trained to always close any open covers.

United Airlines was quick to respond to requests concerning the cameras. United has been on their best behavior following the outcry they heard after the forcible removal of a passenger by security, followed by the death of a puppy that staff insisted be stowed in the overhead compartment.

In speaking about the issue of the seatback cameras to Buzzfeed News, United said that all cameras found in the back of premium seats will now be covered, but that their purpose was never to monitor passengers. Rather, the cameras were included for possible future applications such as video conferencing.

But again… there's just no good way to be sitting in a seat with a camera lens peering back at you.

Cranky Old Guy Tip of the Week

I’m tired of Firefox prompting me about whether I want websites I visit to be able to display desktop web notifications whenever they want to. I can’t think of anything that’s going to be more prone to abuse going forward than these WebNotifications. So I figured that my beloved Firefox browser, amid the bazillion tweakable settings listed under “about:config” would provide that ability… and sure enough!

On the “about:config” page search for: “webnotifications” You’ll get four hits and you want to toggle the first one: “dom.webnotifications.enabled” … toggle it to “false”.

Chrome makes this possible too: Open the settings at the top right, click Settings and scroll down to the bottom where you’ll find “Advanced”. Click it to extend the page into the Advanced settings. In the first section under "Privacy and security," click Site settings. Five entries down find “Notifications”. Flip the switch at the top “Ask before sending (recommended)”. When you turn it off the text will change to “Blocked”.

Test it here:
https://developer.mozilla.org/en-US/docs/Web/API/Notifications_API/Using_the_Notifications_API

Miscellany

Samsung 4G LTE Network Extender 2 for Verizon
https://www.verizonwireless.com/products/samsung-4g-lte-network-extender-2/
$250 -- but no service fee, etc.

- Enhances 4G LTE coverage for homes or businesses up to 7,500 square feet (about a 50–foot radius)
- Supports up to 14 active devices at once

I've mostly recovered from Sunday's Game of Thrones war of the dead vs the living.

From: blog.grc.com
On April 28th, last Sunday, Derek Hamilton posted to blog.grc.com:
Hi Steve, You’ve mentioned many times on SN about the big release event you are going to do on SN. Are we counting weeks, months or longer here? I can’t wait to watch. Thanks.

We're at SQRL Release Candidate #3.

And speaking of SQRL, last week's mess with the just-minted EV code signing certificate appears to be history. Window Defender was first to recognize that SQRL was not malicious and the Windows SmartScreen folks responded quickly to my notifying them that this was from and by a legitimate developer.

Website Credential Stuffing Attacks

Site: https://www.recordedfuture.com/credential-stuffing-attacks/
PDF: https://go.recordedfuture.com/hubfs/reports/cta-2019-0425.pdf

Credential stuffing attacks are Internet-based username and password logon guessing at scale. This requires not only the technology to implement the attacks -- which is now widely available with so many hundreds of thousands of Internet-connected computers available to compromise. But to be profitable these attacks also require the presence of a sufficiently mature marketplace into which a frighteningly large number of discovered working logon credentials can be resold. We've never really talked much about the so-called Dark Web… but it's a real place.

Consequently, the guys are Recorded Future have titled their report and analysis:

"The Economy of Credential Stuffing Attacks"

As their title suggests, they've examined this from the standpoint of the economic incentives which drive this relatively new style of attack.

They begin by explaining:

This report covers the current threat landscape of credential stuffing attacks. It reviews the most popular tools used by cybercriminals to initiate credential stuffing and describes some of the most popular marketplaces that sell compromised credentials. This report contains information gathered using the Recorded Future® Platform, as well as additional open source, dark web, and underground forum research, and will be of most interest to analysts protecting e-commerce, telecommunications, and financial organizations from credential stuffing attacks, as well as those looking for investigative leads on threat actors performing such attacks.

To setup their presentation, they explain...

The rapid proliferation of automated marketplaces on the dark web, fueled by the widespread availability of support infrastructure such as account-checking software, email and password combo lists, and proxy service providers, has created the perfect attack landscape for the abuse of thousands of popular web services such as e-commerce, financial services, travel websites, and telecommunications companies. It is safe to assume that almost every large organization with an online retail presence has had their users exposed to credential stuffing attacks in the past few years, with some companies having upwards of millions of exposed login credentials available for purchase on the dark web at any given moment.

Overall, the main takeaways for this were:

- The first widespread credential stuffing attacks were observed in late 2014, coinciding with the proliferation of automated underground marketplaces. When selling accounts, attackers offered the quick and easy monetization of compromised account credentials. Some actors who engaged in credential stuffing attacks remain active today.

- With an investment of as little as $550, criminals could expect to earn at least 20 times the profit on the sale of compromised login credentials.

- The overall supply of compromised login credentials across several large marketplaces exceeds tens of millions of accounts.

- We have identified at least six popular variants of account-checking software used by cybercriminals; however, dozens of lesser-known variants can be found on the dark web.

- While some companies may choose to implement multi-factor authentication (MFA), which blocks the credential stuffing attack vector, organizations may not be prepared to choose security over convenience.

They write:

Around late 2014 and in the beginning of 2015, we observed the widespread adoption of new dark web business models specifically tailored to facilitate a high volume of trades in a fully automated manner. Designed to emulate legitimate retail platforms such as eBay and Amazon, these so-called "automated shops" allow even low-level criminals to become vendors of stolen data, such as [in this case] compromised login credentials, without having to worry about maintaining their own infrastructure or marketing campaigns. By and large, the adoption of account marketplaces was made possible primarily by the proliferation of account-checking software, or simply "checkers," used as the main tool in credential stuffing attacks.

[I'll note that we have previously heard about the related marketplaces for stolen credit card information. So what we have on the Dark Web is an underground, effectively retail, automated marketplace for the buying and selling of stolen goods. What has occurred recently is the addition of a new class of goods to be sold through this facility: Large quantities of random usernames and passwords for specific websites.]

Compromised account credentials were always a valuable commodity in the dark web — the number of transactions was relatively small, and they were primarily conducted either on a peer-to-peer basis or via semi-automated markets such as AlphaBay, Silk Road, and Hansa Market. In older models, buyers received their wares only after the seller manually approved the deal and delivered the purchased data. Moreover, sellers had to maintain the listings and communicate with the buyers personally.

However, with the advent of automated shops, the need for manual engagement was eliminated and the business of compromised accounts fully transitioned from peer-to-peer dealings to a much more democratized, open-to-everyone enterprise.

For a nominal 10 to 15 percent commission deducted from the amount of each sale, members can upload any number of validated compromised accounts, which in addition to email and password, often include data such as the account holder's city or state of residency, transaction history, and/or account balance. All of this is valuable data to fraudsters seeking to buy accounts tailored to their specific needs. The vendor's main focus is replenishing the stock, while all customer support, remittances, and dispute resolutions are handled by the shop's support team.

At first, only a handful of select vendors became the primary suppliers of stolen data, but as the tradecraft was shared among members of the criminal underground, the business of stolen credentials has grown exponentially.

Since regular internet users tend to reuse the same passwords across multiple websites, threat actors quickly learned that instead of attempting to obtain access to an individual account, which may take a very long time, they should instead focus on hacking multiple random accounts, reducing their efforts.

A combination of several elements made the hacking of various online services accounts not just effortless, but also incredibly lucrative. To launch account brute-forcing, also known as credential stuffing attacks, an attacker only needed… and here are the requirements:

- Brute-forcing software
- A database of random email and password combinations, and
- Access to a pool of proxies.

Early versions of checkers were made to target a single company and were sold for between $50 and $250, depending on the tool's capabilities. These tools would attempt to log in to a website using an email and password combination obtained from a random database often obtained on the dark web. If a combination worked, it would be marked as valid. If not, the software would simply pick another combination from the list and attempt to log in again. For valid logins, more expensive and complex checkers would also collect additional information from the compromised account, such as linked banking and payment card information, account balances, the owner's address, and even transaction history. Until this day, the ingenuity of the method truly lies in the economy of scale, allowing criminals to process hundreds of thousands of combinations in a very short period of time.

Eventually, several dominant players such as STORM, Black Bullet, and Sentry MBA entered the market with more robust tools, supporting an unlimited number of custom plugins, also called "configs," which essentially offered hackers the capability to target almost any company with an online retail presence.

What had initially started as several hundred or several thousand compromised accounts quickly ballooned to hundreds of thousands, or even millions, of accounts. Some of the most prominent account shops have tens of millions of compromised accounts for sale at any given moment. Although the competition quickly brought the average price of a single compromised account from over $10 down to a mere $1 to $2, the overall profitability of credential stuffing attacks increased significantly through sheer volume.

According to underground chatter observed over time, the average success rate for credential stuffing is anywhere between one to three percent. Hence, for every one million random combinations of emails and passwords, attackers can potentially compromise between 10,000 and 30,000 accounts. Moreover, the same database could then be reused over and over again to hack dozens of different websites, yielding even higher profits.

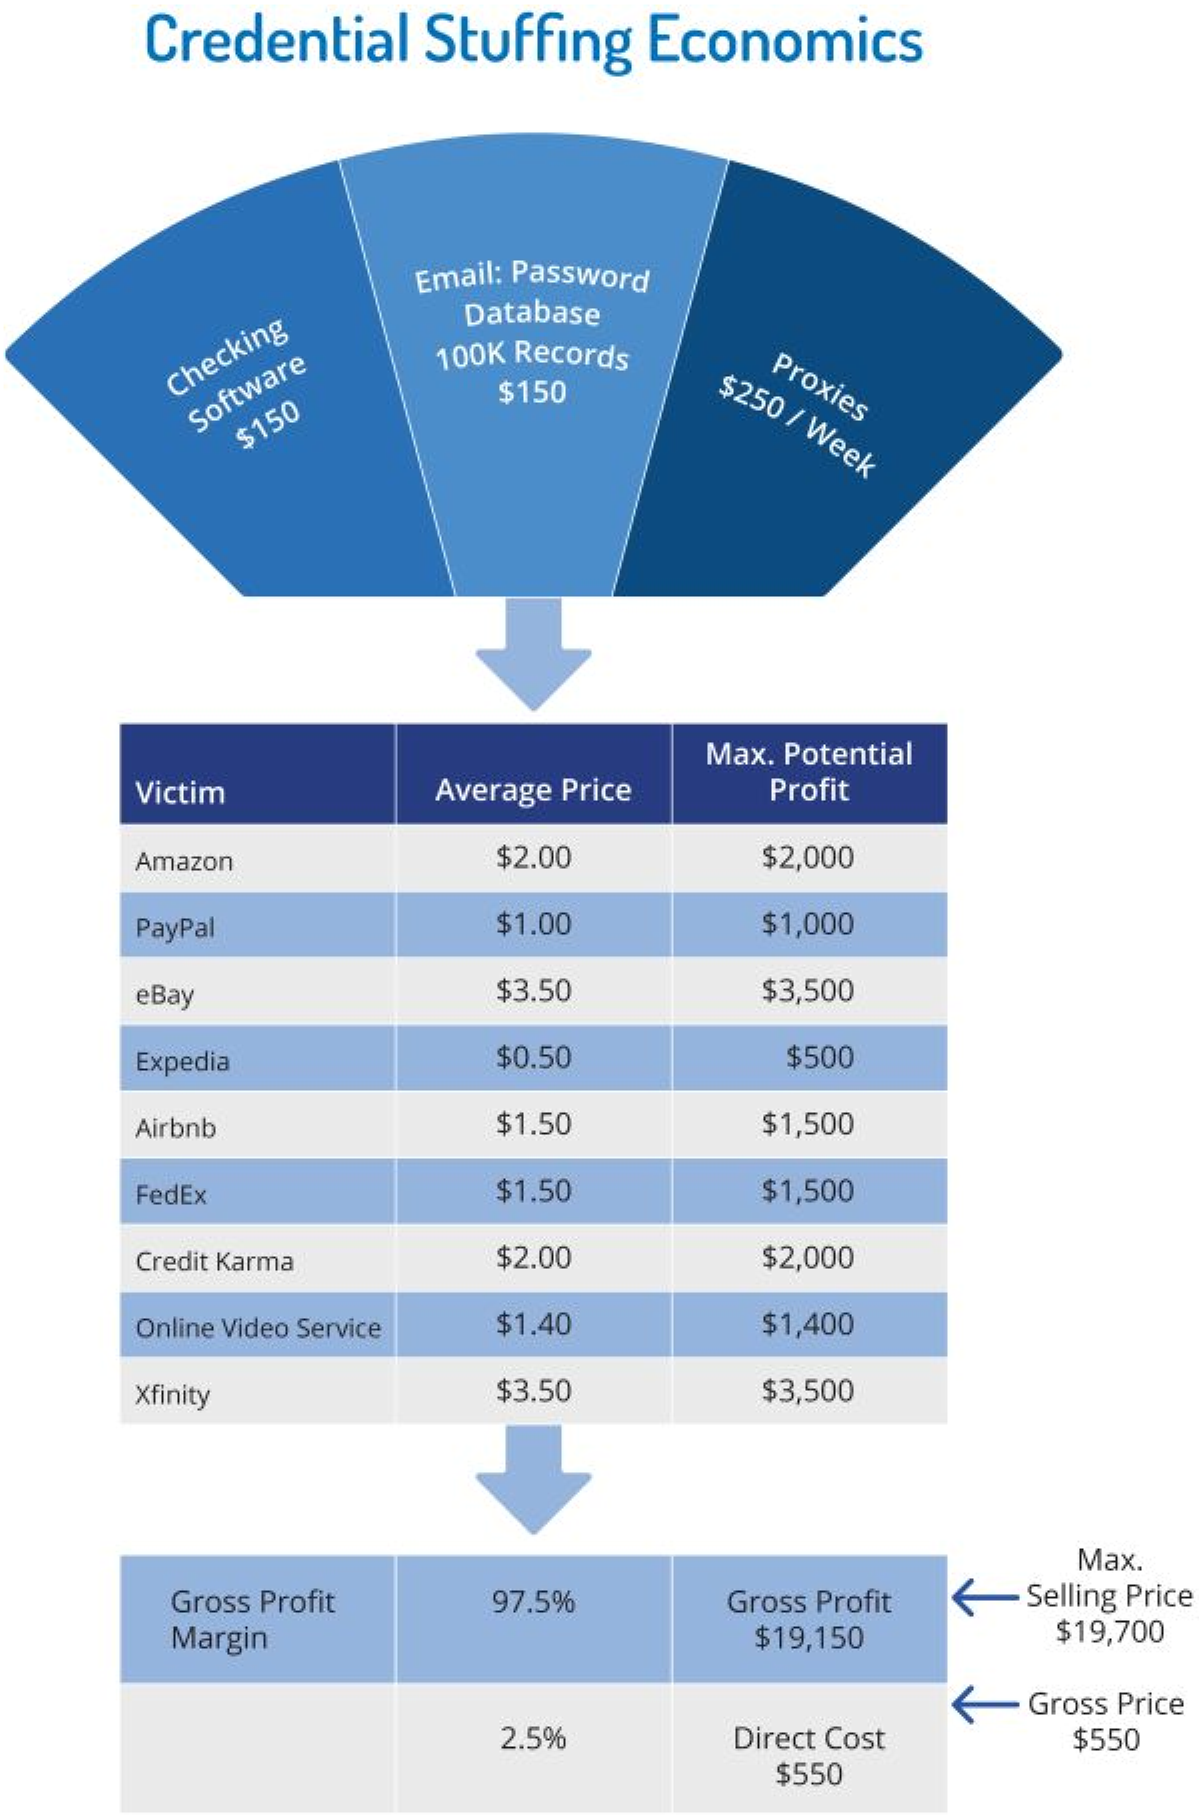

Credential Stuffing Economics

Victim	Average Price	Max. Potential Profit
Amazon	$2.00	$2,000
PayPal	$1.00	$1,000
eBay	$3.50	$3,500
Expedia	$0.50	$500
Airbnb	$1.50	$1,500
FedEx	$1.50	$1,500
Credit Karma	$2.00	$2,000
Online Video Service	$1.40	$1,400
Xfinity	$3.50	$3,500

Gross Profit Margin	97.5%	Gross Profit $19,150	Max. Selling Price $19,700
	2.5%	Direct Cost $550	Gross Price $550

They then go on to examine six specific tools being used to carry out these attacks.

I won’t go into them in detail since anyone who is interested can following the links I have provided at the top of this coverage. My goal was to share an awareness of the existence of this very active underground market and the financial incentives that promise to keep it alive and growing for the foreseeable future.

There are CAPTCHA-bypass features in these tools, though they note that the increasing adoption of multi-factor authentication is helping… at least until websites also begin losing their MFA keys, since MFA also relies upon websites keeping secrets.

~30~