Security Now! #681 - 09-18-18

The Browser Extension Ecosystem

This week on Security Now!

This week we prepare for the first ever Presidential Alert unblockable nationwide text message, we examine Chrome's temporary "www" removal reversal, checkout Comodo's somewhat unsavory marketing, discuss a forthcoming solution to BGP hijacking, examine California's forthcoming IoT legislation, deal with the return of Cold Boot attacks, choose not to click on a link that promptly crashes any Safari OS, congratulate Twitter on adding some auditing, check in on the Mirai Botnet's steady evolution, look at the past year's explosion in DDoS number of size, note another new annoyance brought to us by Windows 10... Then we take a look at the state of the quietly evolving web browser extension ecosystem.

How many web browsers have how many browser extensions installed?

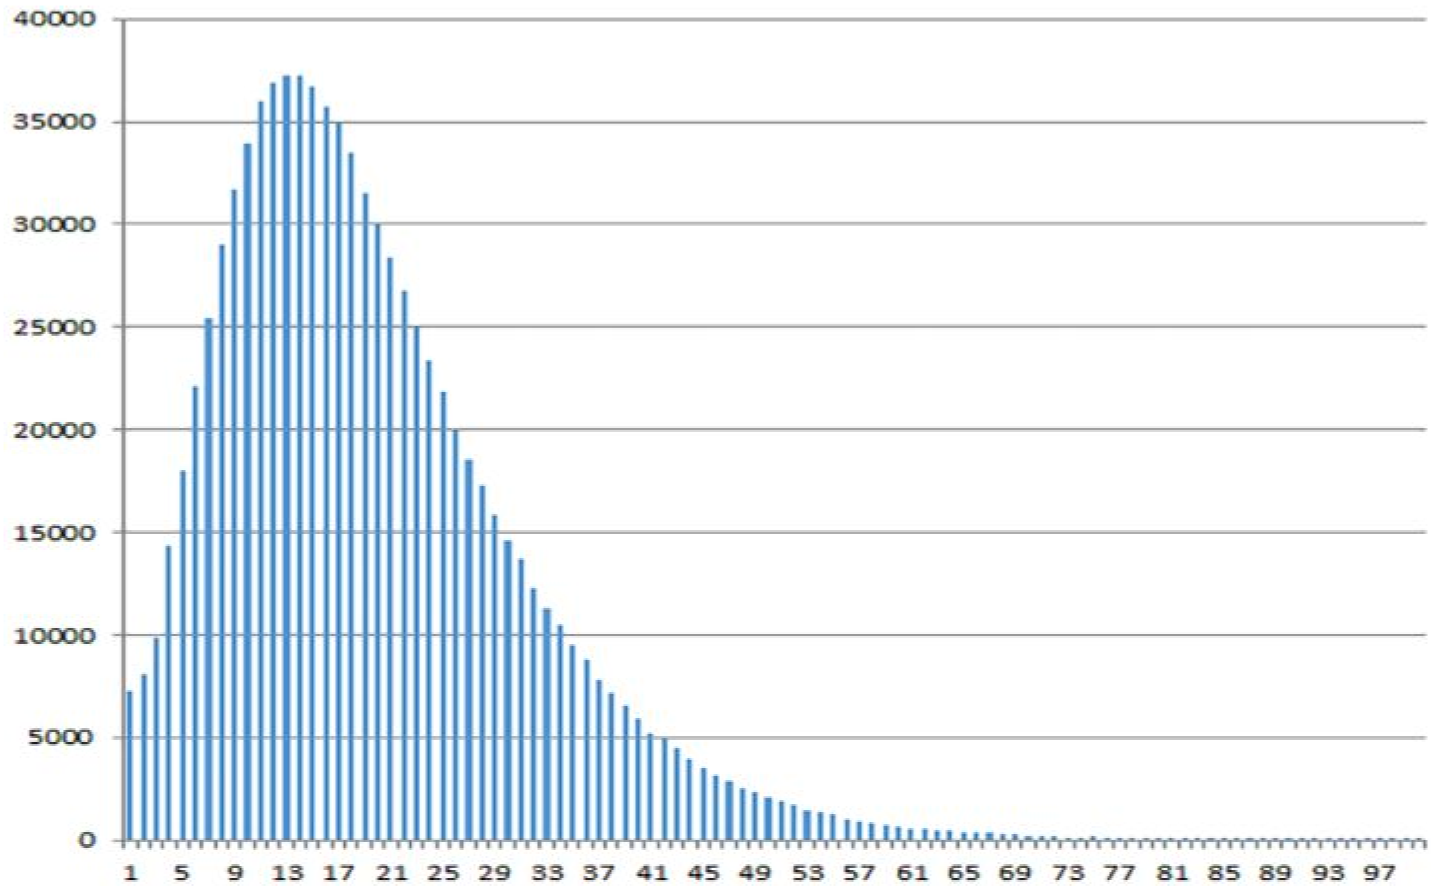

Security News

Postponed: FEMA's first "Presidential Alert" emergency test message to EVERYONE.
https://www.fema.gov/emergency-alert-test

The National EAS and WEA test has been postponed to the backup date of October 3, 2018, beginning at 2:18 p.m. EDT.

The first ever EAS (Emergency Alert System) and WEA (Wireless Emergency Alerts) test was originally scheduled for this coming Thursday. That date has now been bumped back to Wednesday, October 3rd, two weeks and one day from today due to ongoing response to the recent hurricane Florence.

FEMA's announcement states that: The WEA portion of the test commences at 2:18 p.m. EDT, and the EAS portion follows at 2:20 p.m. EDT. The test will assess the operational readiness of the infrastructure for distribution of a national message and determine whether improvements are needed.

Although this is will be the forth Emergency Alert System test, it will be the first test of the new Wireless Emergency Alert test. Previous EAS tests were conducted in November of 2011, September of 2016 and September of 2017.

So... on October 3rd, all cell towers will begin broadcasting the WEA test, for approximately 30 minutes, beginning at 2:18 p.m. EDT. During this time, WEA compatible cell phones that are switched on, within range of an active cell tower, and whose wireless provider participates in WEA should be capable of receiving the test message. Some cell phones will not receive the test message, and cell phones should only receive the message once. The WEA test message will have a header that reads "Presidential Alert" and text that says:

“THIS IS A TEST of the National Wireless Emergency Alert System. No action is needed.”

FEMA said that more than 100 carriers will be participating in the test.

Chrome 69… to WWW or not to WWW? That is the question!
WWW is back, baby! (Kinda…)

From the Chromium Blog:

> In Chrome M69, we rolled out a change to hide special-case subdomains “www” and “m” in the Chrome omnibox. After receiving community feedback about these changes, we have decided to roll back these changes in M69 on Chrome for Desktop and Android.
>
> In M70, we plan to re-ship an adjusted version: we will elide “www” but not “m.” We are not going to elide “m” in M70 because we found large sites that have a user-controlled “m” subdomain. There is more community consensus that sites should not allow the “www” subdomain to be user controlled.

> We plan to initiate a public standardization discussion with the appropriate standards bodies to explicitly reserve “www” or “m” as special case subdomains under-the-hood. We do not plan to standardize how browsers should treat these special cases in their UI. We plan to revisit the “m” subdomain at a later date, after having an opportunity to discuss further with the community. Please leave feedback, or let us know about sites that have user-controlled subdomains at “www” or “m” in this bug.

On top of this, as I briefly noted last week, Chrome's initial implementation was also buggy and would match and remove multiple "www's" from a single domain name... not just a leading "www". This they have fixed.

However, users who are voicing opinions remain unhappy with the Chromium team's stated intention to return to removing a leading "www" from URL in Chrome 70. The Chromium team solicited comments and they received many. Here are three samples:

Comment 1:

> What is the reasoning behind the decision not to wait for the standardization discussions? As the majority market share holder, it's incumbent on Chrome and Chromium to be cautious when making arbitrary decisions about how connection information should be displayed- or in this case, deliberately transformed and obscured. The urgency with which this change is being pushed is baffling.
>
> I strongly disagree with the idea that we should be connecting to one host name while displaying another in the address bar. Absent legitimate discussions with the appropriate standards bodies, this decision feels myopic at best, and wrong-headed at worst.

Comment 2:

> Standardization discussions aside, no changes to the look and function of sub-domains should happen, unless this is an opt-in setting. SaaS, ESP and ISP providers take great care to manage their domains and use 301 redirects to bring users to the right location. A browser should not determine this function unless it is completely a user choice.
>
> Perhaps Google/Chromium would like to respond as to the exact end game and reason for such a change. Dictating a change which is controlled through DNS, but shows up differently on a browser is just plain confusing and wrong. I don't agree with what Safari and Windows have done either.

Comment 3:

> At least make it a toggleable option, I don’t mind if URLs are elided by default. But please have an option to disable this.

Comodo's latest marketing campaign:

Sean Nelson (@seannelson)

> I manage a couple hundred School District websites that use Let's Encrypt certs for HTTPS. LE renews automatically using a script on my servers. I just got a breathless phone call from Comodo warning me that one of my domains has a cert that will expire soon (like 70 days from now) and I better buy a 3 year certificate or I would be sorry. When I told her that I know what I'm doing, and it renews automatically, she made it sound like I was being irresponsible by risking the certificate renewal every 3 months rather than every 3 years.
>
> I wonder how many other LE users are getting these phone calls from Comodo desperately trying to scare them into buying a certificate instead of using a free LE cert.

It's an interesting marketing approach. Since every website's TLS certificate clearly displays its expiration data, desperate sales agents with nothing better to do could attempt to roust users whose certs are nearing retirement and renewal.

Thwarting BGP Hijacking by creating a new system of trust

The final piece of the standard to protect against BGP hijack attacks gets first official draft
An large effort led by the NIST and DHS has been working on a solution to BGP troubles.

The effort is termed SIDR for Secure Inter-Domain Routing and consists of three interrelated protocols, the first two of which were finalized a year ago, last September.

RFC8206: BGPsec Considerations for Autonomous System (AS) Migration
This document discusses considerations and methods for supporting and securing a common method for Autonomous System (AS) migration within the BGPsec protocol.

RFC8210: The Resource Public Key Infrastructure (RPKI) to Router Protocol, Version 1
In order to verifiably validate the origin Autonomous Systems and Autonomous System Paths of BGP announcements, routers need a simple but reliable mechanism to receive Resource Public Key Infrastructure (RFC 6480) prefix origin data and router keys from a trusted cache. This document describes a protocol to deliver them.

And now, a year later, we have the third and final piece of the solution with the BGP Route Origin Validation (ROV) standard which, when coupled with the other two existing protocols, promises to help ISPs and cloud providers protect against BGP hijack attacks.

https://www.nccoe.nist.gov/sites/default/files/library/sp1800/sidr-piir-nist-sp1800-14-draft.pdf
(A mind-numbing 264 page document.)

Executive Summary:

- It is difficult to overstate the importance of the internet to modern business and to society in general. The internet is essential to the exchange of all manner of information, including transactional data, marketing and advertising information, remote access to services, entertainment, and much more.
- The internet is not a single network, but rather a complex grid of independent interconnected networks. The design of the internet is based on a trust relationship between these networks and relies on a protocol known as the Border Gateway Protocol (BGP) to route traffic among the various networks worldwide. BGP is the protocol that internet service providers (ISPs) and enterprises use to exchange route information between them.

- Unfortunately, BGP was not designed with security in mind. Traffic typically traverses multiple networks to get from its source to its destination. Networks implicitly trust the BGP information that they receive from each other, making BGP vulnerable to route hijacks.

- A route hijack attack can deny access to internet services, misdeliver traffic to malicious endpoints, and cause routing instability. A technique known as BGP route origin validation (ROV) is designed to protect against route hijacking.

- The National Cybersecurity Center of Excellence (NCCoE) at the National Institute of Standards and Technology (NIST) has developed proof-of-concept demonstrations of BGP ROV implementation designed to improve the security of the internet's routing infrastructure.

- This NIST Cybersecurity Practice Guide demonstrates how networks can protect BGP routes from vulnerability to route hijacks by using available security protocols, products, and tools to perform BGP ROV to reduce route hijacking threats. The example implementation described in this guide aims to protect the integrity and improve the resiliency of internet traffic exchange by verifying the source of the route.

BENEFITS

The NCCoE's practice guide is intended to improve the security and stability of the global internet by allowing networks to verify the validity of BGP routing information and strengthen the security and stability of traffic flowing across the global internet benefitting all organizations and individuals that use and rely on it. This practice guide can help your organization:

- reduce the number of internet outages due to BGP route hijacks

- ensure that internet traffic reaches its destination

- make informed decisions regarding routes and what actions to take in cases when BGP ROV implementation has not been performed or has indicated that an advertised route is invalid

California approves some baby step IoT legislation.
https://www.zdnet.com/article/first-iot-security-bill-reaches-governors-desk-in-california/
https://nakedsecurity.sophos.com/2018/09/13/california-bill-regulates-iot-for-first-time-in-us/

The California State legislature recently approved ‘SB-327 Information privacy: connected devices’ and handed it to governor Jerry Brown to sign it into law. The legislation introduces security requirements for connected devices sold in the US. It defines these devices as any device that connects directly or indirectly to the internet and has an IP or Bluetooth address... which is kinda everything.

The legislation reads:

<quote> *“This bill, taking effect on January 1, 2020, will require a manufacturer of a connected device, as those terms are defined, to equip the device with a reasonable security feature or features that are appropriate to the nature and function of the device, appropriate to the information it may collect, contain, or transmit, and designed to protect the device and any information contained therein from unauthorized access, destruction, use, modification, or disclosure, as specified.”*

Of course, that says a whole lot of absolutely nothing.

However, the bill also says:

<quote> *“If a connected device is equipped with a means for authentication outside a local area network the authentication system must meet one of two criteria”*:

- *If the device uses a default password, the password must be unique to each device; or,*

- *The device must prompt users to set up their own password whenever the user sets up the device for the first time.*

Last Monday, in reaction to this, Errata Security's Rob Graham generated a long posting from which I’ll share the first three paragraphs:
https://blog.erratasec.com/2018/09/californias-bad-iot-law.html

California has passed an IoT security bill, awaiting the governor's signature/veto. It’s a typically bad bill based on a superficial understanding of cybersecurity/hacking that will do little improve security, while doing a lot to impose costs and harm innovation.

It’s based on the misconception of adding security features. It’s like dieting, where people insist you should eat more kale, which does little to address the problem you are pigging out on potato chips. The key to dieting is not eating more but eating less. The same is true of cybersecurity, where the point is not to add “security features” but to remove “insecure features”. For IoT devices, that means removing listening ports and cross-site/injection issues in web management. Adding features is typical “magic pill” or “silver bullet” thinking that we spend much of our time in infosec fighting against.

We don’t want arbitrary features like firewall and anti-virus added to these products. It’ll just increase the attack surface making things worse. The one possible exception to this is “patchability”: some IoT devices can’t be patched, and that is a problem. But even here, it’s complicated. Even if IoT devices are patchable in theory there is no guarantee vendors will supply such patches, or worse, that users will apply them. Users overwhelmingly forget about devices once they are installed. These devices aren’t like phones/laptops which notify users about patching.

Cold Boot Attacks are back
Ten years ago, in 2008, we talked about the problem of sleeping PCs and laptops which shutdown all power but kept RAM alive... // Freezing, etc.

But it turned out that just rebooting a sleeping machine would not wipe its RAM and if a hostile OS got control it could often find all kinds of goodies left behind in RAM... including the master keys to the system's whole-disk encryption.

That was ten years ago and BIASes were updated to proactively wipe RAM and to often be more reluctant to be booted from externally supplied media.

However, last Thursday two researchers from F-Secure gave a presentation at the Sec-T Conference titled: "An ice-cold Boot to break BitLocker". The talk's teaser read:

<quote> *A decade ago, academic researchers demonstrated how computer memory remanence could be used to defeat popular disk encryption systems. Not much has happened since, and most seem to believe that these attacks are too impractical for real world use. Even Microsoft has started to play down the threat of memory remanence attacks against BitLocker, using words such as “they are not possible using published techniques”.*

Well... We will publish techniques that allow recovery of BitLocker encryption keys from RAM on most, if not all, currently available devices. While BitLocker is called out in the title, the same attacks are also valid and fully effective against other platforms and operating systems.

So what did these guys do?

Keeping the target machine -- a laptop in this case -- powered up, they opened the back hooked a readily available FLASH programmer to the machine's BIOS setting storage chip, then tweaked some settings to disable the RAM memory overwrite and enable booting from external devices.

They then carried out a traditional cold boot attack by booting a special program from a USB stick.

And since this defeats BitLocker or other disk encryption, not only the contents of RAM, but also anything stored on the disk is accessible externally without the protection layers provided by the OS, since the OS has been bypassed.

What to do?

Maintain physical security for any sleeping machines.

As with the recent TPM bypass involving the ACPI sleep state, if possible disable the sleep option in the BIOS and also in the machine's OS. In Windows it's possible to do this with the Group Policy Editor:

Run "gpedit.msc" from the RUN... prompt.

Navigate to: Computer Configuration > Policies > Administrative Templates > System > Power Management. Then...

Disable the use of Sleep mode with these group policies:

- Sleep Settings > 'Allow Standby States (S1 – S3) when sleeping (on battery)' – DISABLED
- Sleep Settings > 'Allow Standby States (S1 – S3) when sleeping (plugged in)' - DISABLED

(Play with hibernation settings there, too, if you wish.)

Change the power button and lid close actions as follows:

- Button Settings > 'Select the lid switch action (on battery)' – HIBERNATE
- Button Settings > 'Select the lid switch action (plugged in)' - HIBERNATE
- Button Settings > 'Select the Power button action (on battery)' - HIBERNATE
- Button Settings > 'Select the Power button action (plugged in)' - HIBERNATE

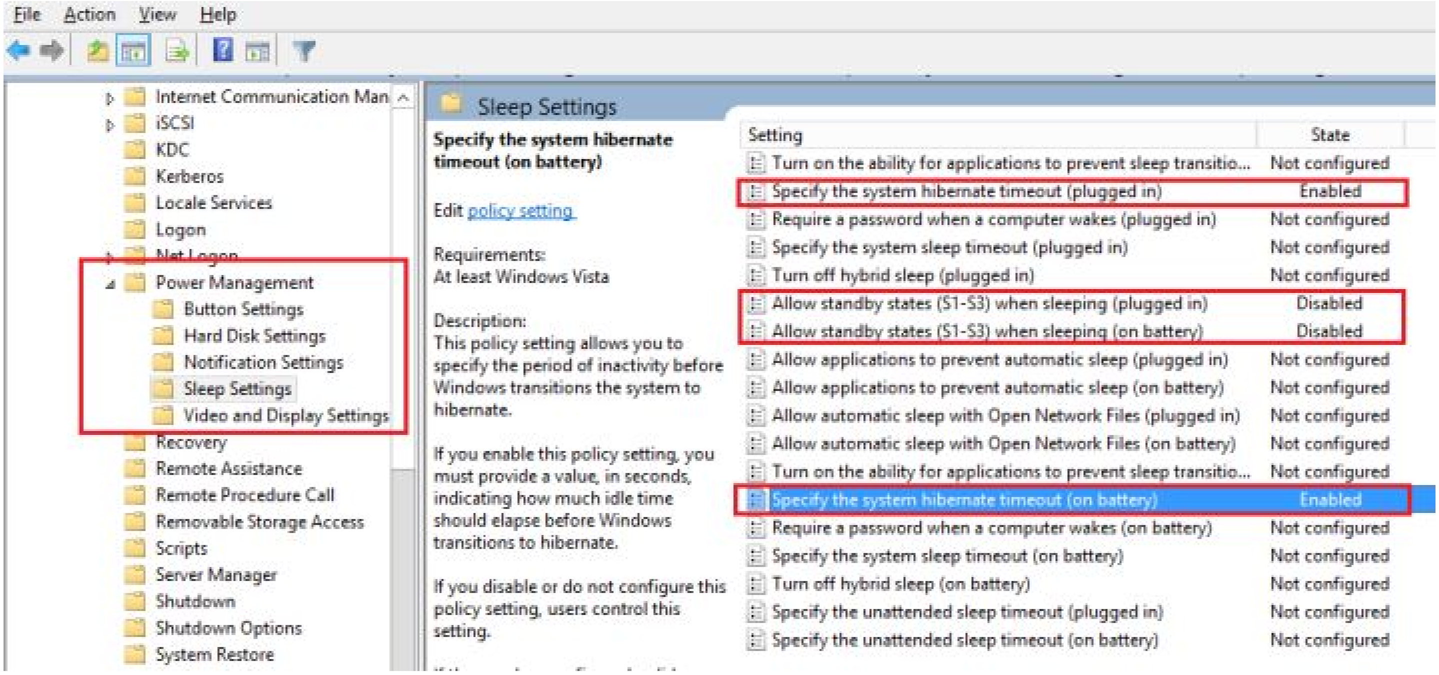

Specify the system hibernate timeout (on battery)

Edit policy setting

Requirements:
At least Windows Vista

Description:
This policy setting allows you to specify the period of inactivity before Windows transitions the system to hibernate.

If you enable this policy setting, you must provide a value, in seconds, indicating how much idle time should elapse before Windows transitions to hibernate.

If you disable or do not configure this policy setting, users control this setting.

Setting	State
Turn on the ability for applications to prevent sleep transitio...	Not configured
Specify the system hibernate timeout (plugged in)	Enabled
Require a password when a computer wakes (plugged in)	Not configured
Specify the system sleep timeout (plugged in)	Not configured
Turn off hybrid sleep (plugged in)	Not configured
Allow standby states (S1-S3) when sleeping (plugged in)	Disabled
Allow standby states (S1-S3) when sleeping (on battery)	Disabled
Allow applications to prevent automatic sleep (plugged in)	Not configured
Allow applications to prevent automatic sleep (on battery)	Not configured
Allow automatic sleep with Open Network Files (plugged in)	Not configured
Allow automatic sleep with Open Network Files (on battery)	Not configured
Turn on the ability for applications to prevent sleep transitio...	Not configured
Specify the system hibernate timeout (on battery)	Enabled
Require a password when a computer wakes (on battery)	Not configured
Specify the system sleep timeout (on battery)	Not configured
Turn off hybrid sleep (on battery)	Not configured
Specify the unattended sleep timeout (plugged in)	Not configured
Specify the unattended sleep timeout (on battery)	Not configured

Safari-Ripper: An Apple Safari vulnerability that quickly crashes iOS and macOS devices
A relatively new and processor-intensive feature of CSS, known as "Backdrop Filter" uses 3D acceleration to process the image background behind CSS elements.

Backdrop-filter blurs or color shifts the region behind an element. It is a heavy processing task, and some software engineers and web developers have speculated that the rendering of this effect takes a toll on iOS' graphics processing library, eventually leading to a crash of the mobile OS altogether.

Security research, Sabri Haddouche was poking around while researching reliable denial of service (DoS) bugs against various browsers when he discovered a simple way to brain Safari-based systems to their knees.

Sabri simply used nested divs with that property to quickly consume all graphic resources and freeze or kernel panic the OS.

And the problem could be even more widespread since Apple forces all browsers and HTML-capable apps listed on the App Store to use its WebKit rendering engine -- which is where the trouble lies. So this flaw will probably crash any app that's capable of loading a web page.

<<< This link =WILL= crash any vulnerable browser: >>>
https://cdn.rawgit.com/pwnsdx/ce64de2760996a6c432f06d612e33aea/raw/23f2faa0aadb4babbfd228c8bb32a26a8c51c741/safari-ripper.html

The link DID crash my fully updated pre-iOS v12 iPad... =AND= my just updated iOS v12 iPad.

Twitter has added permitted "Apps and devices" management
Goto "Settings and Privacy" then "Apps and devices"
https://twitter.com/settings/sessions

- First it's promoting a connection to Facebook.
- Then: "Apps connected to your Twitter account"
- Then: "Recently used devices to access Twitter"
- Three devices showed as last logged on December 31st, 1969. Huh?

The Mirai botnet is on track to become historic.
The total number of exploits in the latest version’s repertoire has grown to 16. And while the majority remain aimed at compromising routers, network video recorders and DVRs... Mirai is now also targeting unpatched Apache Struts vulnerabilities.

A close cousin of Mirai, dubbed Gafgyt, is targeting a known vulnerability in SonicWall devices.

Researchers believe that both Botnet agents are being sourced by the same threat actor since both samples were hosted at the same domain.

Palo Alto Network's Unit 42, who have been tracking these 'Bots have suggested that "The incorporation of exploits targeting Apache Struts and SonicWall by these IoT/Linux botnets could be an indication of a larger movement from consumer device targets to enterprise targets."

The explosion in DDoS attack sizes
https://www.nexusguard.com/threat-report-q2-2018
The DDoS mitigation company "NexusGuard" (not a business I would like to be in!) has published a report comparing DDoS attacks during the first two quarters of 2017 versus the first two quarters of 2018.

During the identical periods, 2018 saw a 29% upwards jump in the number of attacks and a staggering 543% increase in average attack size with that new average now being 26.37 Gigabits per second of attacking bandwidth.

The NexusGuard's report shows that the average size of attacks in Q2 2017 was 4.10 Gbps with a maximum observed of 63.70 Gbps.

One year later, for Q2 2018, this average size increased over 500% to 26 Gbps and a maximum size attack was seen of 359 Gbps.

The attacking sources will come as no surprise to Security Now! followers. NexusGuard wrote:

The increase in attacks and their sizes is being attributed to attackers amassing giant botnets using insecure IoT devices. Attackers are using vulnerabilities in these devices to rapidly build large botnets that can then be used to perform targeted attacks that are increasingly difficult to stop.

For example, at one point the Mirai Satori botnet was seen from over 280,000 IP addresses over a 12 hour period and the newer Anarchy botnet was able to amass over 18,000 routers in a single day. These botnets were created by attackers exploiting vulnerabilities in routers such as ones made by Huawei & D-Link.

NexusGuard wrote: *"In addition, severe botnet epidemics like last year’s Satori continued to threaten cyberspace by exploiting zero-day vulnerabilities. Since its high-profile attack on Huawei home routers in December 2017, Satori has wreaked havoc over the past few months on various IoT devices, including: GPON-capable routers manufactured by South Korea’s Dasan, D-Link’s DIR-620 routers, and XiongMau uc-httpd 1.0.0 IoT devices. Additionally, the quarter saw the emergence of the Anarchy botnet, which exploited zero-day vulnerabilities in a similar fashion as Satori."*

Slightly more than half of attacks (55.28%) had a duration of less than 90 minutes, but the average attack lasted 318.10 minutes long. The average was pushed up because some attacks lasted for many days, with the longest being 6 days, 5 hours, and 22 minutes.

Although nearly 23/rds of attacks (64.13%) were under 10 Gbps, the average size was 26.37 Gbps.

The United States was the largest source of attacks at 20%, followed by China, France, Germany, and Russia.

Windows 10 starting to "warn" users who attempt to install another browser
(You really can't make this stuff up.)

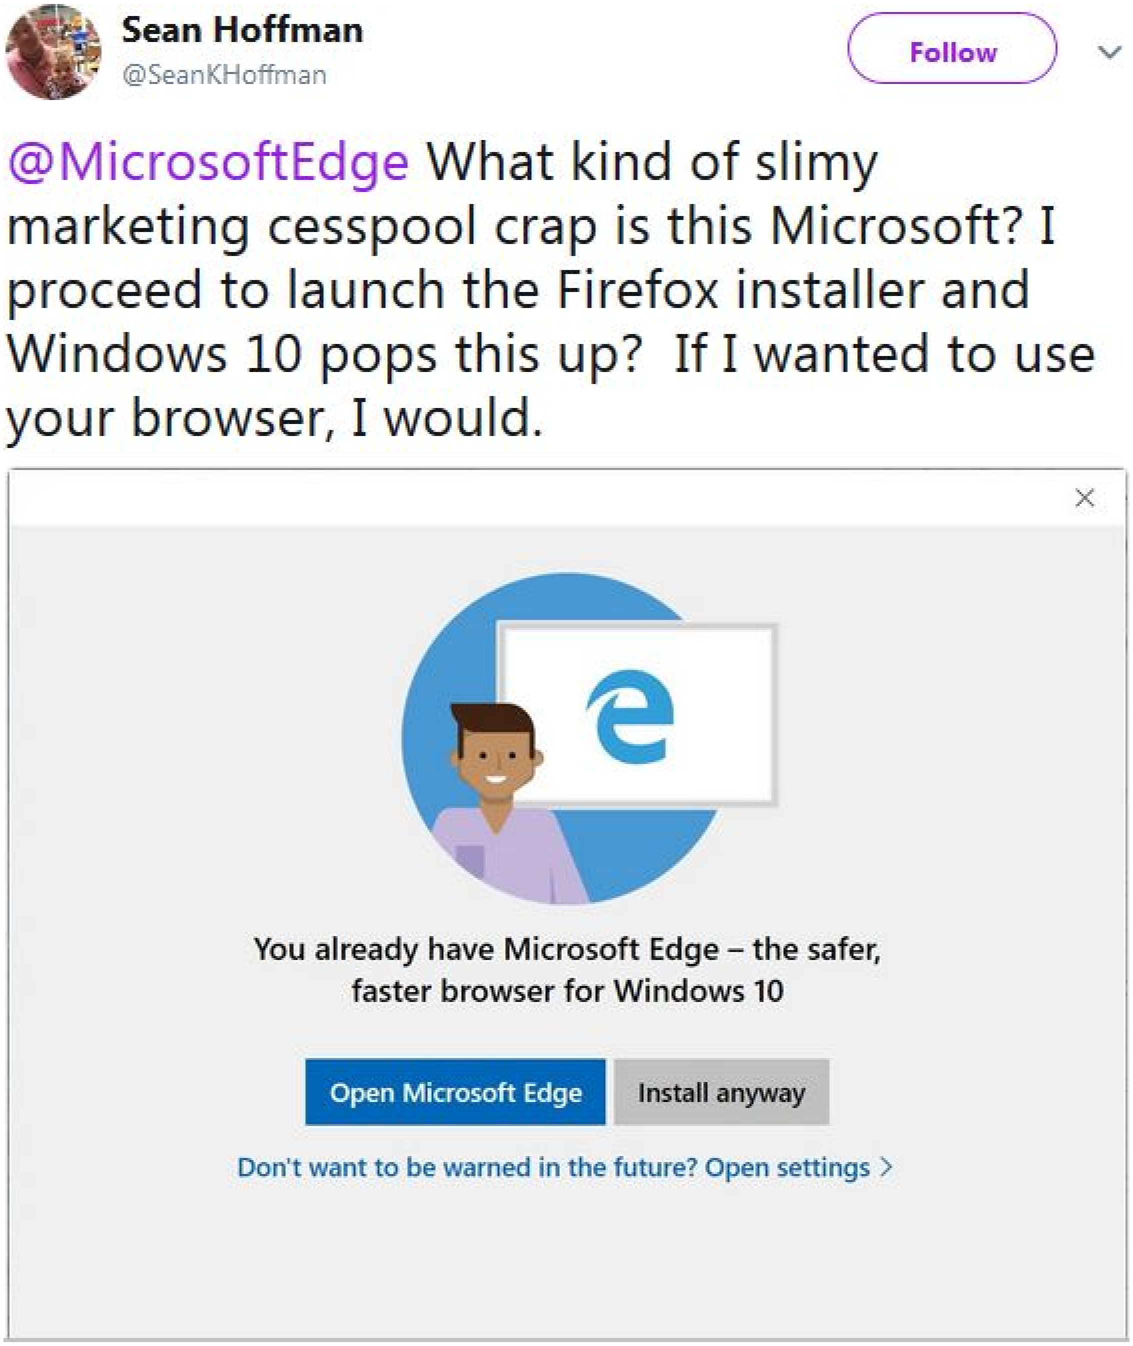

Sean Hoffman
@SeanKHoffman

@MicrosoftEdge What kind of slimy marketing cesspool crap is this Microsoft? I proceed to launch the Firefox installer and Windows 10 pops this up? If I wanted to use your browser, I would.

You already have Microsoft Edge – the safer, faster browser for Windows 10

Don't want to be warned in the future? Open settings >

Microsoft has started to test a new feature in Windows 10 Insider builds that displays a message promoting Microsoft edge when a user attempts to install an alternative web browser such as Firefox, Chrome, Vivaldi, or Opera.

This is currently the Windows 10 Insiders builds only, but this is where many similar idea are born.

Note that Google is similarly naggy when using their search from a non-Chrome browser. I have to keep saying "No thanks" or just ignoring the dialog. But at least it's not blocking what I want to do until I confirm that I'm happy with what I already have.

Errata

Alexey in Russia
Subject: Couple of thoughts in defense of MikroTik (and they *can* autoupdate ;-)
Date: 18 Sep 2018 04:28:40

Hi Steve and Leo!

Long-time listener, SpinRite licensee with all the customary bla-bla-blas to you both :)

I've been using MikroTik (MT for short) routers since about 2012, I have a couple at home and one at work (I'm a sysadmin in a small company). I'd like to mention that my oldest router still receives ALL the updates (generally, one-two a month) – it's a very strong point for MT in my book. I've heard about Ubiquity long after becoming acquainted with MT, and never had reason to switch nor spare cash to experiment, despite your fascination with infamous Ubiquity Edge.

First of all I'd like to mention that MT routers can easily be configured to autoupdate. Details are below, if you are interested.

Secondly, I *do* agree with you MT *have* to change their policy and default settings. And they have the capabilities to do so without additional hassle for anyone: they have several templates for user to choose from on initial setup, they just have to harden their “Home router” template and make it the default for new devices. They don’t. But if they did that users would be reasonably protected, and advanced ones will configure from scratch or choose another template.

However in their defense I have to say that (A) it's *absolutely clear* from even a cursory look, that their routers are not consumer-oriented plug-and-pray - or plug-and-forget - devices and (B) that they are targeted not at casual home users but at geeks at the very least (I'll leave out professionals and such here).

In my experience configuring MT router is a bit more like configuring Linux server than configuring a “traditional” consumer router. Common good practice to use non-standard ports for services -including ssh, Web UI and Winbox if you use them – fully applies here. And to disable ones you don't use as well. And you absolutely have to configure your firewall – default configuration is far too basic for deployment.

There's no checkbox to allow/disable WAN access, we're grownups here, you have to configure firewall rules – with netfilter/iptables-like cli commands or in a GUI – to allow or disallow such connections.

And yes, with great power comes great responsibility: MT routers can easily be misconfigured.

I have to say: all recently reported vulnerabilities you told us about were mitigated by properly configured firewall (i.e. allow inputs only from trusted IPs in LAN or maybe WAN if you really need it, drop all other inputs) and changing default service ports helps a lot too.

=== Regarding autoupdate ===

As you've said in your coverage RouterOS (OS of MT's routers) have extensive scripting capabilities and cron-like scheduler. Script to check for update, install it automatically and reboot is about four or five lines long. Couple more if you want to be notified by email before or instead of updating. Schedule it to run daily at 3 or 4 o'clock in the morning and you're good to go.

Judging from MT forum posts, sysadmins and ISP guys don't think their routers should autoupdate. They prefer to test new releases beforehand and only deploy security fixes, let's say. But for home users who don't want to do it manually it's there) And user don't need to buy a new router, their old one are perfectly capable of doing so)

Me personally – I prefer to do my updates manually. And now I have an automatic update available notification)

Miscellany

Jeff Root in San Diego, CA
Subject: SN-680: OpenVPN Running with Privilege?
Date: 15 Sep 2018 20:17:22
:
All the reporting I saw on this bug said it was exploitable on Windows only.

Checking my Linux system, I see that OpenVPN is installed (by default on Debian, at least) so that it runs as user "nobody" and group "nogroup".

The config file for OpenVPN clearly states that it will drop privileges after initialization, _except_ on Windows.

As far as I'm concerned, the question that should be asked is: why can't it drop privs on Windows? Is that not supported by the Windows API? Which should, actually, be the fix for this anyway.

This violates the fundamental security principle of "least privilege", which has been known and implemented since the 1960's.

SpinRite

Justin in Olympia, WA
Subject: SpinRite Catches the Bad Guy
Date: 16 Sep 2018 00:34:43
:
Steve,

Long time SN listener, blah blah blah.

Wanted to let you know how SpinRite helped to catch a criminal. Granted it was in a roundabout way, but still - it probably wouldn't have happened without SpinRite.

My security camera server was having issues. The normally rock solid software started crashing on a frequent basis. Everytime it happened, until I happened to notice it and was able to reboot the server, my system was down. It was getting really annoying.

I'd tried a lot of things to fix it, but I was getting the sneaking suspicion that the SSD start up drive might be the problem. Earlier this week, I ran SpinRite on level 2. While it didn't find any bad sectors, it did the trick. The camera software hasn't crashed since.

This is a VERY good thing, as today someone tried to steal some things I had in front of my house. I was *praying* that the camera software hadn't crashed and was thus able to record the attempt. And thanks to SpinRite the entire incident was recorded.

I was able to provide the clips of the guy to our Sheriff's office, and within an hour, he was in custody.

Oh yeah, he fought with the Deputy when he was being arrested, and got introduced to the business end of the K-9 unit. He also had several felony warrants outstanding - A real upstanding member of society.

Thanks to SpinRite a really bad dude is behind bars.

Thanks again!

The Web Browser Extension Ecosystem

I was launched into taking a closer look at the current state of our web browser extension ecosystem by the news that two major versions from now -- Firefox 64 -- will be adding a welcome UI feature: The pop-up menu of its web browser extensions currently lists commands such as "Manage the Extension", "Remove from Toolbar", "Pin to Overflow Menu", and also non-extension related options to Display or Hide the Menu Bar or the Bookmarks Bar or further customize the browser's UI. Firefox 64 will add a prominent "Remove Extension" at the top of the pop-up menu list.

The team at Mozilla wants to make the process of removing unwanted extensions easier and very very clear. At some point in the future, though apparently not with FF64, they further plan to build-in an option to submit an abuse report on the spot to solicit feedback about the reason for the extension’s removal.

This news, along with our recent coverage of malicious extensions creeping into Chrome and the many times I have encountered weird and apparently unwanted browser toolbars installed in other people's machines led me to look into the current ecosystem of web browser extensions.

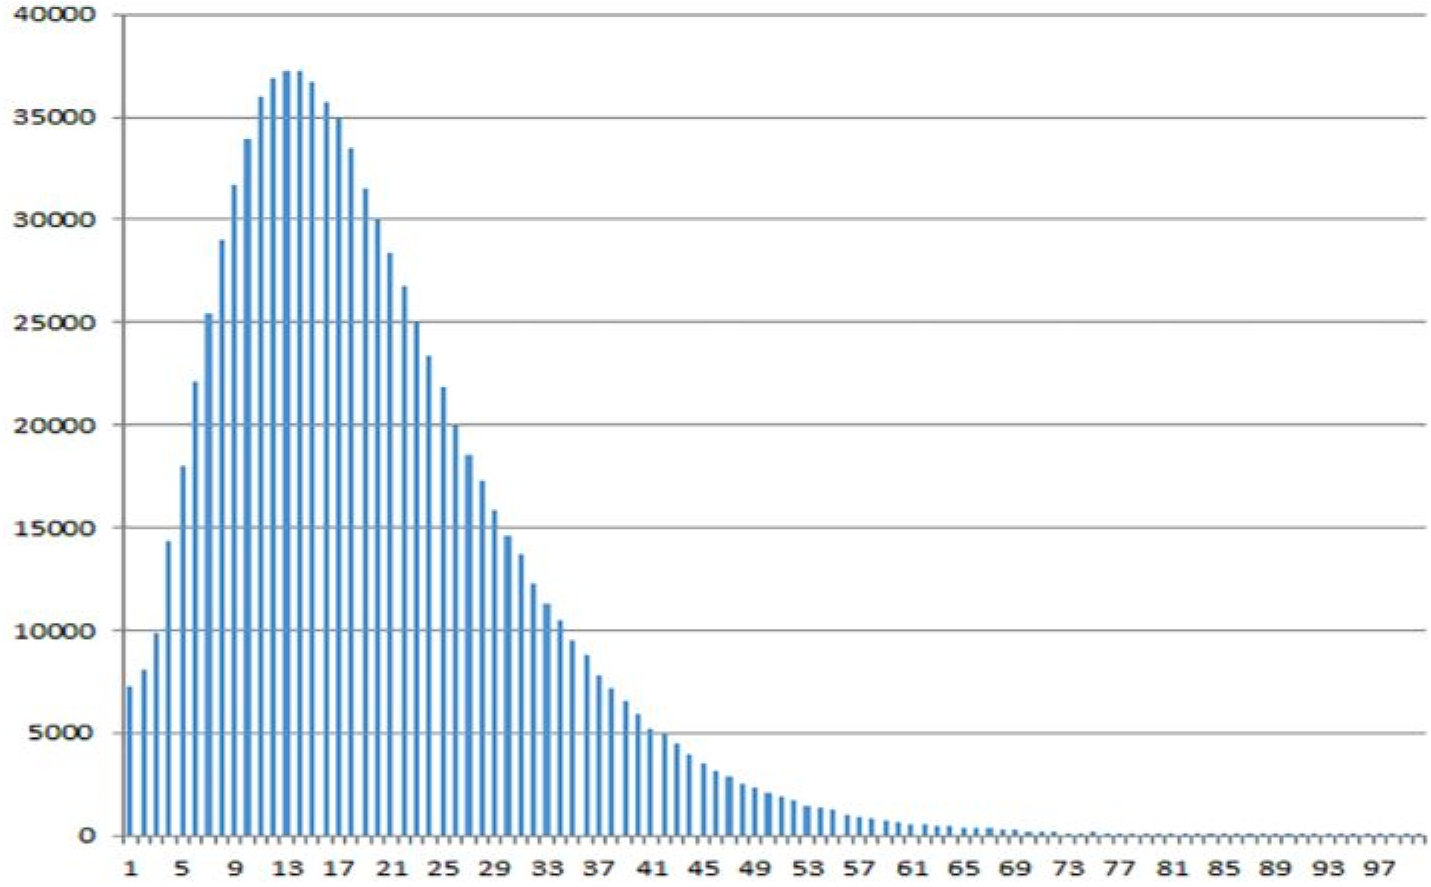

We've been covering the welcome advances being made in performance, security, and usability with both Google and Firefox. Browser extensions (like other forms of software clogging our machines) tend to accumulate over time and can eventually have a significant impact upon browser performance... and security.

The privileges of the inherently untrusted JavaScript loaded by aggressive pages we visit are actively constrained. (Thank goodness!) But, for browser extensions to be able to extend fundamental browser functions to do the things we want, those extensions need to be given an inherently greater level of trust and access.

Mozilla and Google have reputations that allow us to trust them with the fundamental operation of their products. But, for the most part, and as we've recently covered, browser extensions sourced from unknown parties can turn out not to have the best interests of their users in mind. Browser extensions are free or voluntarily user supported. So the underlying motivations of the extension's author are not always clear. Reputation is all we have to judge by, and most extensions are from unknown sources.

https://journals.plos.org/plosone/article?id=10.1371/journal.pone.0179281
https://journals.plos.org/plosone/article/file?id=10.1371/journal.pone.0179281&type=printable

In digging around I uncovered an interesting and recent research paper from the summer of 2017 titled: "Quantifying the web browser ecosystem"

Abstract:

Contrary to the assumption that web browsers are designed to support the user, an examination of a 900,000 distinct PCs shows that web browsers comprise a complex ecosystem with millions of addons collaborating and competing with each other. It is possible for addons to “sneak in” through third party installations or to get “kicked out” by their competitors without user involvement. This study examines that ecosystem quantitatively by constructing a large-scale graph with nodes corresponding to users, addons, and words (terms) that [self-]describe addon functionality. Analyzing addon interactions at user level using the Personalized PageRank (PPR) random walk measure shows that the graph demonstrates ecological resilience. Adapting the PPR model to analyzing the browser ecosystem at the level of addon manufacturer, the study shows that some addon companies are in symbiosis [whereas] others clash with each other as shown by analyzing the behavior of 18 prominent addon manufacturers. Results may herald insight on how other evolving internet ecosystems may behave, and suggest a methodology for measuring this behavior. Specifically, applying such a methodology could transform the addon market.

Introduction:

Web browsers have become a major component of the routine human-computer interaction, with some operating systems based entirely on browsers (e.g., ChromeOS by Google [1]). Browser extensions, also known as addons, are computer programs that (as the name suggests) extend, improve, and personalize browser capabilities. More than 750 million addons were downloaded and installed by Google Chrome browser users as of June 2012 [2].

Some examples of addons include an extension that allows visually impaired users to access the content of bar charts on the Web [3], an extension that addresses users' security concerns by seamlessly producing a unique password for each website the user accesses [4].

Internet software companies are very interested in installing their addons, and particularly toolbars, on users' machines. Toolbars are GUI widgets that typically reside in the upper part of the browser's window, extending the browser's functionality. Toolbars can collect information about the browsing history of the user (e.g., Yahoo! Toolbar [5]) and can redirect user search activity to a specific search portal (e.g. MyWebSearch.com). Crucially, the company that owns the search portal, and typically also the toolbar, receives payments from ad providers per user click on the ads it displays (primary ad providers are Google and Yahoo!). This revenue generation model is used extensively by software companies that distribute freeware products [6]. For example, 45% of AVG Antivirus Technologies sales in 2012 were generated by its browser toolbar [7]. It was estimated that Google, the biggest Web advertising firm, might have lost 1.3 billion in revenue in 2013 because of changes to its policy with respect to toolbars and a resulting shift of some addon distributors to Google's competitors [8, 9].

Consequently, addons compete with each other over resources (such as battery, memory, disk space, and computing power) and user attention. Regardless of how intelligent they are, they may be aware of each other and may "piggyback" on each other, or uninstall each other. Addon behavior within the Web browser is characterized by addons making their own decisions independently and often unbeknown to the user, which comprises a complex ecosystem with the user being just one of the participants. A key issue in understanding that ecosystem, responding to it, regulating it, and transforming it into a mature market is the current inability to show that it is inherently stable and measurable. This study addresses that issue.

So... I think our takeaway from this is the recognition that the end user (and also apparently the browser's vendor) are in less control over this browser extension world that has evolved than they might think... and that those browser extensions we may have installed some time ago, and then promptly forgotten, are still alive and installed and having access to everything we do on the web... unencrypted.

So I welcome tools which profile our extensions and alert us when an extension is slowing down the browser's launch or page loads, and anything our web browsers can do to provide us with additional decision making tools and empowerment will be welcome.

Our operating systems have pretty much locked down the management of what gets installed into our machines, and when... but the browser extension ecosystem hasn't yet matured to that level.

Just as it's a good idea to occasionally curate the apps that have access to our Google, Facebook and Twitter accounts, taking an occasional survey of our web browser add-ons see what's there, or what might have crawled in uninvited, makes good security and privacy sense.

~30~