Security Now! #685 - 10-16-18

Good Samaritans?

This week on Security Now!

This week we observe the untimely death of Microsoft's co-founder Paul Allen, revisit the controversial Bloomberg China supply chain hacking report, catch up on Microsoft's October patching fiasco, follow-up on Facebook's privacy breach, look at the end of TLS v1.0 and 1.1, explore Google's addition of control flow integrity to Android 9, look at a GAO report about the state of US DOD weapons cybersecurity, consider the EOL of PHP 5.x chain, take a quick look at an A/V comparison test, entertain a few bits of feedback from our listeners, and then consider the implications of grey-hat vigilante hacking of others' routers.

Security News

Bloomburg / China Implant Followup
Business Insider:
The security community increasingly thinks a bombshell Bloomberg report on Chinese chip hacking could be bogus
https://www.businessinsider.com/security-community-voicing-increasing-doubts-about-bombshell-bloomberg-chinese-chip-hacking-2018-10

Business Insider reports: "All parties involved have denied the report, including, most recently, secretary of the Department of Homeland Security, during a Senate hearing."

Last Wednesday, the secretary of the Department of Homeland Security denied the report in a Senate hearing — the strongest on-the-record government denial yet. Kirstjen Nielsen said: "With respect to the article, we at DHS do not have any evidence that supports the article. We have no reason to doubt what the companies have said."

But shortly after saying she had no evidence in favor of the story, DHS secretary Nielsen said: "What we can tell you though, is it's a very real and emerging threat that we're worried about."

And a manufacturing expert told Business Insider: "I don't actually think it's hard to inject stuff that the brand or design team didn't intentionally ask for."

Then, during the same hearing, the FBI's Director Chris Wray said that he couldn't confirm nor deny the existence of any investigation into compromised SuperMicro equipment.
Why couldn't he simply deny it?

TechCrunch's Zack Whittaker, who has been reporting on security for years, considers this confusion of press releases and statements to be the nature of this kind of reporting. When the stakes are this high, reporting on security vulnerabilities and classified information means that you're more often than not dealing with making your sources anonymous to protect them, which, rightfully, opens your work up to denials and condemnations from the organizations you cover. In TechCrunch, Zack wrote:

The sources are anonymous — likely because the information they shared wasn't theirs to share or it was classified, putting sources in risk of legal jeopardy. But that makes accountability difficult. No reporter wants to say "a source familiar with the matter" because it weakens the story. It's the reason reporters will tag names to spokespeople or officials so that it holds the powers accountable for their words. And, the denials from the companies themselves — though transparently published in full by Bloomberg — are not bulletproof in outright rejection of the story's claims. These statements go through legal counsel and are subject to government regulation. These statements become a counterbalance — turning the story from an evidence-based report into a "he said, she said" situation.

Recall, for example, that even here in the U.S., companies who are legally forced to divulge information to law enforcement are also bound to deny that any such disclosure ever took place. Whatever it was that happened, occurred several months ago, so there's been ample time for

plenty of behind-the-scenes court orders for anyone involved.

If this DID happen, then those who were most affected know that it happened and there's no way they won't be taking steps to prevent it in the future.

And even if it did NOT happen, no one -- not a single expert -- has said that it could not happen. All experts have acknowledged the possibility. Which means that the upshot of this scare -- if that's all it was -- is that the possibility is now more on everyone's mind than before. And that's a good thing.

And Kevin Mitnick just demonstrated a malicious USB cable entering keystrokes into a fully patched Windows 10 machine and installed malware. Recall that a virtual USB keyboard was the way the Yubico system entered the password for the user -- by emulating a connected USB keyboard. And Windows has a wide array of handy hotkey shortcuts for doing all sorts of things.

So... the threat itself is real. And if we have seen anything, it's that what CAN be done is eventually done.

https://www.bleepingcomputer.com/news/technology/facebook-states-30-million-people-affected-by-last-months-view-as-bug/

Facebook downgrades the severity of the "View As" pivot attack. Only 30 million account were affected.
October 12, 2018 / An Update on the Security Issue
https://newsroom.fb.com/news/2018/10/update-on-security-issue/

As Lawrence Abram's wrote for his BleepingComputer site when covering this news…
https://www.bleepingcomputer.com/news/technology/facebook-states-30-million-people-affected-by-last-months-view-as-bug/

First, quoting Facebook's update: "We now know that fewer people were impacted than we

originally thought. Of the 50 million people whose access tokens we believed were affected, about 30 million actually had their tokens stolen."

Lawrence: "Isn't that great? Only 30 million!"

Facebook:
First, the attackers already controlled a set of accounts, which were connected to Facebook friends. They used an automated technique to move from account to account so they could steal the access tokens of those friends, and for friends of those friends, and so on, totaling about 400,000 people. In the process, however, this technique automatically loaded those accounts' Facebook profiles, mirroring what these 400,000 people would have seen when looking at their own profiles. That includes posts on their timelines, their lists of friends, Groups they are members of, and the names of recent Messenger conversations. Message content was not available to the attackers, with one exception. If a person in this group was a Page admin whose Page had received a message from someone on Facebook, the content of that message was available to the attackers.

The attackers used a portion of these 400,000 people's lists of friends to steal access tokens for about 30 million people. For 15 million people, attackers accessed two sets of information – name and contact details (phone number, email, or both, depending on what people had on their profiles). For 14 million people, the attackers accessed the same two sets of information, as well as other details people had on their profiles. This included username, gender, locale/language, relationship status, religion, hometown, self-reported current city, birthdate, device types used to access Facebook, education, work, the last 10 places they checked into or were tagged in, website, people or Pages they follow, and the 15 most recent searches. For 1 million people, the attackers did not access any information.

I'll reiterate something that's made clear from this, but which I haven't seen noted anywhere else: This was a fully targetable attack -- and both of the top two Facebook executives were compromised. To my thinking, while 30 million is a big number, and that number clearly means that automation was used to pivot from one compromised user to the next, the fact that View As can show how =any= specified Facebook user would view one's account, and the bug then enabled a pivot over to that now-authenticated user's identity, meant that anyone on Facebook could have at least some of their information compromised.

And in this update, Facebook appears to have neglected to mention that the owner of someone's authenticated Facebook identity could then use the widespread "Sign in with Facebook" to logon to many other Internet websites and services under the identity of that impersonated account.

Facebook added:
People can check whether they were affected by visiting our Help Center. In the coming days, we'll send customized messages to the 30 million people affected to explain what information the attackers might have accessed, as well as steps they can take to help protect themselves, including from suspicious emails, text messages, or calls.

https://www.facebook.com/help/securitynotice?ref=sec

Microsoft's Windows 10 October Mess

This month's Win10 update/patch/upgrade debacle has been such a mess that it was a bit

challenging trying to organize it to create a coherent view for this podcast.

Last Tuesday afternoon, while we were recording last week's podcast, Microsoft posted: "Updated version of Windows 10 October 2018 Update released to Windows Insiders: https://blogs.windows.com/windowsexperience/2018/10/09/updated-version-of-windows-10-october-2018-update-released-to-windows-insiders/

<quote> Last week we paused the rollout of the Windows 10 October 2018 Update (version 1809) for all users as we investigated isolated reports of users missing files after updating. Given the serious nature of any data loss, we took the added precaution of pulling all 1809 media across all channels, including Windows Server 2019 and IoT equivalents. We intentionally start each feature update rollout slowly, closely monitoring feedback before offering the update more broadly. In this case the update was only available to those who manually clicked on "check for updates" in Windows settings. At just two days into the rollout when we paused, the number of customers taking the October 2018 Update was limited. While the reports of actual data loss are few (one one-hundredth of one percent of version 1809 installs), any data loss is serious.

Yes... and what Microsoft failed to note, here, was the sobering news that there HAD previously been MANY reports from Microsoft's own Win10 insiders of this mass data deletion occuring to them... which Microsoft had ignored.

<quote> Prior to re-releasing the October 2018 Update our engineering investigation determined that a very small number of users lost files during the October 2018 Update. This occurred if Known Folder Redirection (KFR) had been previously enabled, but files remain in the original "old" folder location vs being moved to the new, redirected location. KFR is the process of redirecting the known folders of Windows including Desktop, Documents, Pictures, Screenshots, Videos, Camera Roll, etc. from the default folder location, c:\users\username\<folder name>, to a new folder location. In previous feedback from the Windows 10 April 2018 Update, users with KFR reported an extra, empty copy of Known Folders on their device. Based on feedback from users, we introduced code in the October 2018 Update to remove these empty, duplicate known folders. That change, combined with another change to the update construction sequence, resulted in the deletion of the original "old" folder locations and their content, leaving only the new "active" folder intact.

Microsoft has now found and fixed three scenarios where this was seen to occur:

- Using KFR the user redirected a known folder to a different drive. For example, suppose you ran out of space on your C drive. You want to save some files separate from your primary folder, so you add another drive to your system for these. You create "D:\documents" and change the location of the files known folder from the original "old" location c:\users\username\documents to D:\documents. In some cases, if the contents of c:\users\username\documents were not moved to D:\documents, then a user could also encounter this issue. When the October 2018 Update was installed the original "old" folder was deleted including the files in that folder (in this example c:\users\username\documents would be deleted; d:\documents, the new location, would be preserved).
- The user configured one or more of their Known Folders (Desktop, Documents, Pictures, Screenshots, Videos, Camera Roll, etc.) to be redirected (KFR) to another folder on OneDrive.

For example, the user changed the location property of the documents folder from c:\users\username\documents to another folder. During this process the system prompts the user and asks if they would like to move the files to the new location. If the files were not moved and the October 2018 Update is installed the original "old" folder was deleted including the files in that folder.

- The user used an early version of the OneDrive client and used the OneDrive settings to turn on the Auto save feature. This feature turned on KFR for the Documents and/or Pictures folders based on the user's choice but did not move the existing files from the original "old" location to the new location. For example, if a user turned on Auto Save for pictures the location of the Pictures folder would be changed from c:\users\username\pictures to c:\users\username\onedrive\pictures, but no files would be moved. The current version of this feature moves the files. If the files were not moved and the October 2018 Update was installed the original "old" folder was deleted including the files in that folder (in this example c:\users\username\pictures would be deleted; c:\users\username\onedrive\pictures, the new location, would be preserved).

And indirectly addressing the fact that insiders had been screaming about this well before the update's release, and being ignored, Microsoft concluded their posting with:

<quote> To help us better detect issues like this, today we have enabled a new feature in the Windows Insider Feedback Hub. We have added an ability for users to also provide an indication of impact and severity when filing User Initiated Feedback. We expect this will allow us to better monitor the most impactful issues even when feedback volume is low.

We will continue to closely monitor the update and all related feedback and diagnostic data from our Windows Insider community with the utmost vigilance. Once we have confirmation that there is no further impact we will move towards an official re-release of the Windows 10 October 2018 Update. We apologize for any impact these issues may have had on any of our customers. We are committed to learning from this experience and improving our processes and notification systems to help ensure our customers have a positive experience with our update process.

Also...

A bad HP keyboard driver made its way into both the 1803 and 1809 updates and was causing a BSOD on boot for HP machines.
"HP devices may experience blue screen error WDF_VIOLATION after installing HP keyboard driver (version 11.0.3.1)"
https://support.microsoft.com/en-us/help/4468372/hp-devices-may-experience-blue-screen-error-wdf-violation-after-instal

"This driver has a known incompatibility with certain HP devices on Windows 10 versions 1803 and 1809"

User's who had not restarted were able to remove the driver before rebooting. But once the reboot update had instantiated the driver into the system, the system was no longer bootable.. and recovering from that was a mess. (The knowledge base article linked about provides the details.)

Also...
"What needs your attention": Intel audio display notification
https://support.microsoft.com/en-us/help/4465877/what-needs-your-attention-intel-audio-display-notification
<quote> Microsoft and Intel have identified a compatibility issue with a range of Intel Display Audio device drivers that may result in excessive processor demand and reduced battery life. As a result, the update process to Windows 10 October 2018 Update (version 1809) will fail. If you see a "What needs your attention" notification when you run the October Update, you have an Intel Display Audio device driver (intcdaud.sys, versions 10.25.0.3 – 10.25.0.8) installed in your system.

Also... (However)
Some users reported that their audio had died after the 1803 or 1809 updates. The trouble was with the "Intel SST Audio Controller" (Intel Smart Sound Technology driver version 09.21.00.3755). It was installed in error, killing the audio of users. It can be deleted under Device Manager and audio will be restored.

Also...
There appears to be a display brightness resetting problem which is affecting some users who report that their system's display brightness resets to its lowest level after every reboot. No further word on that yet.

Also...
With the release of Windows 10 1809, the Microsoft Edge web browser and Microsoft's UWP Store apps may no longer be able to connect to the Internet... which is a pesky problem for a web browser. it turns out that Edge and UWP Store apps now require TCP/IPv6 to be enabled or they will not be able to connect. Presumably, if someone had previously manually disabled IPv6, those apps would now be broken. So IPv6 needs to be re-enabled.

Given this wide array of problems, is it any wonder that enterprise IT are reluctant to jump onto updates immediately.

And the idea of setting your Win10 machines to holding off on installing new feature and even monthly updates for at least a week or two makes a lot of sense.

Also...

Recall that way back on May 8th of this year, Trend Micro, through their 0-Day initiative, notified Microsoft of a potentially serious remote code execution (RCE) bug in their age-old and present everywhere JET database engine. Microsoft immediately confirmed that they had reproduced the issue, then sat on the bug while Trend Micro gave them 120 days to fix it.

Microsoft said that the fix wasn't going to make it into the SEPTEMBER patches, but the clock had ticked past 120 days, so Trend Micro went public.

This vulnerability is being disclosed publicly without a patch in accordance with the ZDI 120 day deadline.

- 05/08/18 - ZDI reported the vulnerability to the vendor and the vendor acknowledged the report
- 05/14/18 - The vendor replied that they successfully reproduced the issue ZDI reported
- 09/09/18 - The vendor reported an issue with the fix and that the fix might not make the September release
- 09/10/18 - ZDI cautioned potential 0-day
- 09/11/18 - The vendor confirmed the fix did not make the build
- 09/12/18 - ZDI confirmed to the vendor the intention to 0-day on 09/20/18

There was that "MicroPatch" stop-gap measure offered by the Acros Security "0Patch" folks: https://0patch.com/

But given that this vulnerability would be closed with the October update, it didn't seem worth bothering with, especially from a 3rd-party.

The good news was, this remote code execution problem was finally fixed last week after FIVE months:

CVE-2018-8423 | Microsoft JET Database Engine Remote Code Execution Vulnerability
Security Vulnerability
https://portal.msrc.microsoft.com/en-US/security-guidance/advisory/CVE-2018-8423

<quote> A remote code execution vulnerability exists in the Microsoft JET Database Engine.

An attacker who successfully exploited this vulnerability could take control of an affected system. An attacker could then install programs; view, change, or delete data; or create new accounts with full user rights. Users whose accounts are configured to have fewer user rights on the system could be less impacted than users who operate with administrative user rights.

To exploit the vulnerability, a user must open/import a specially crafted Microsoft JET Database Engine file. In an email attack scenario, an attacker could exploit the vulnerability by sending a specially crafted file to the user, and then convince the user to open the file.

Whoops!!! But wait! Or was it?

Patching, Re-Patching and Meta-Patching the Jet Database Engine RCE (CVE-2018-8423)
https://blog.0patch.com/2018/10/patching-re-patching-and-meta-patching.html

<quote> This is a story about a Windows vulnerability that was reported to Microsoft, published as "0day" before the official patch was available, micropatched by us one day later, subsequently patched by Microsoft, found to be incorrectly patched, and now micropatched again over the flawed official patch. Confused? It's actually a nice case study to demonstrate... not how we are smarter than Microsoft (we really aren't) but how micropatching is a much more suitable technology for fixing vulnerabilities on billions of computers than the current de-facto standard of what we call "fat updates".

At this time, so far as we know, this vulnerability is not yet being widely exploited, or exploited in the wild. But if bad guys were gearing up to use it during its brief window of availability, they have just been gifted with at least another month of window.

The Micropatch is an intriguing idea, since the micro-patch itself is only 18 bytes. But the installer is much larger. And you need to "sign up for a free account."

Proof-of-concept code published for Microsoft Edge remote code execution bug
The PoC can be hosted on any website and requires that users press the Enter key just once.
https://www.zdnet.com/article/proof-of-concept-code-published-for-microsoft-edge-remote-code-execution-bug/

Arguing against the idea of waiting too long to apply security patches is this tidbit:

One thing that Microsoft DID patch this month was a potentially potent and trivial to exploit, website deliverable remote code execution bug which, prior applying this months patches (if you

dare!) would allow any site visited to execute arbitrary code on the visitor's computer if they can be induced to press their "Enter" key.

A Kuwaiti security researcher, Abdulrahman Al-Qabandi, after making sure Microsoft had patched this problem last week, which he had previously discovered and responsibly reported via Trend Micro's Zero-Day Initiative, published full in-depth details about the Edge vulnerability on his blog... including the trivial HTML and JavaScript full proof-of-concept details.

We know that "ease of exploitation" matters a lot. Not all of the gremlins inhabiting the Internet are high-end exploit developers. And in this instance, time is not on the attacker's side since the problem with Edge has been resolved. But this exploit would take mere minutes to weaponize, Edge is the default browser for Win10, and until machines are updated they will be vulnerable. it also sounds as though this could be delivered through an advertisement.

October 15th: The Deprecation of TLS v1.0 and v1.1

Deprecation of Legacy TLS 1.0 and 1.1 Versions
https://webkit.org/blog/8462/deprecation-of-legacy-tls-1-0-and-1-1-versions/

Modernizing Transport Security
https://security.googleblog.com/2018/10/modernizing-transport-security.html

Modernizing TLS connections in Microsoft Edge and Internet Explorer 11
https://blogs.windows.com/msedgedev/2018/10/15/modernizing-tls-edge-ie11/

Removing Old Versions of TLS
https://blog.mozilla.org/security/2018/10/15/removing-old-versions-of-tls/

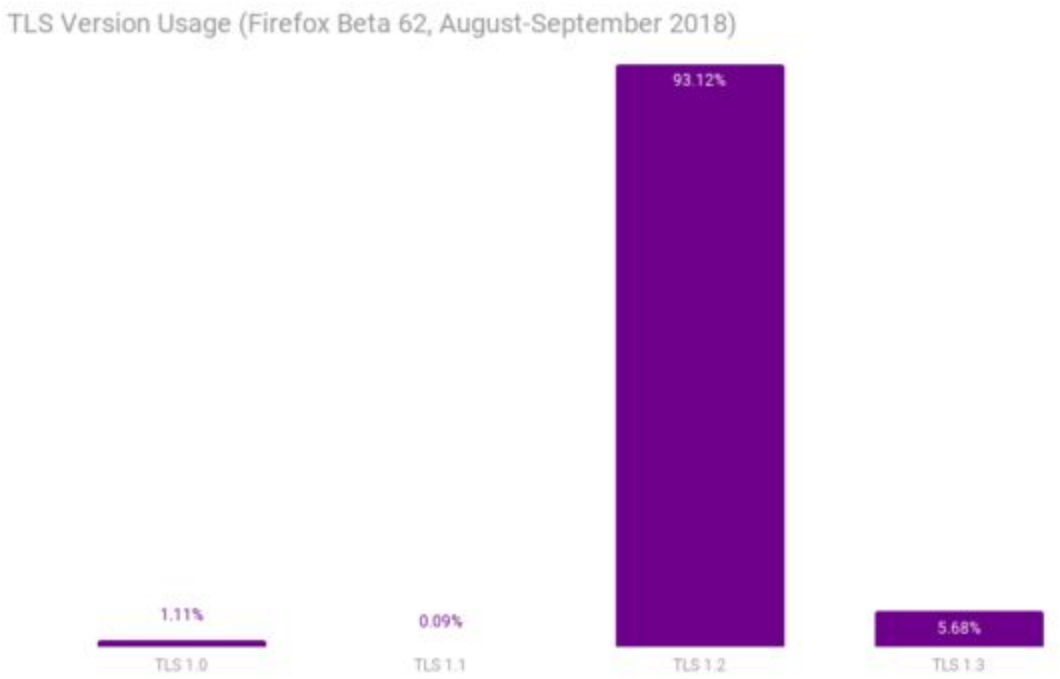

TLS versions for all connections established by Firefox Beta 62, August-September 2018

TLS v1.0 and v1.1 are formally put to sleep on March 2020

TLS v1.0 (which was mostly a renaming of SSL v3.0) will enjoy its 20th birthday this coming January 2019. As such, and by any measure, it has been a spectacularly successful Internet protocol which, even today, nearly 20 years after its introduction, is without any flaws serious enough to have forced its early retirement. But… nevertheless, its age is beginning to show through its lack of support for 20 years of subsequent progress.

Consequently, yesterday, October 15th, all four of the major browser vendors -- Apple with Webkit, Google with Chrome, Microsoft with Edge and IE and Mozilla with Firefox -- coordinated their announcements of the end-of-life for both TLS v1.0 and v1.1.

It's time to move on. All of the browser vendors see usage of TLS v1.0 down near 1% with v1.1 never having really even replaced v1.0, so it's less than 1/10th of 1%. At the same time, v1.2 is enjoying more than 93% usage and the next generation TLS v1.3 is beginning to gain connection share with nearly 6%.

As usual, the deprecation of both v1.0 and v1.1 will be a staged and made clear. But websites beware: Unless you're able to accept and terminate connections over v1.2 by march of 2020, the Internet will suddenly become a very lonely place.

Google's Chrome will begin deprecating v1.0 and v1.1 in Chrome 72. Visitors to sites that are unable to accept v1.2 or v1.3 connections will begin to see deprecation warnings in the DevTools console in that release. Then v1.0 and v1.1 will be disabled altogether in Chrome 81. This will affect users on early release channels starting January 2020.

Google Adds Control-Flow Integrity to Beef up Android Kernel Security

CFI is an outgrowth of the LLVM project work. The "LLVM" project began as a research project at the University of Illinois where its goal was to explore the design and implementation of a modern, static single assignment computer language compilation system capable of supporting both static and dynamic compilation of arbitrary programming languages.

Today it is a state-of-the-art collection of modular and reusable compiler and toolchain technologies which are often preferred to the older GCC tools and toolchain.

While it's tempting to assume that LLVM is an acronym for Low Level Virtual Machine, as things turned out the project has little to do with traditional virtual machines so that the name "LLVM" is not an acronym for anything. However, the project does work with an "Intermediate Representation" (IR) which is, in a sense a representation of a low level virtual machines.

In compiler design, static single assignment (SSA) is a property of an intermediate representation (IR), which requires that each variable is assigned exactly once, and every variable is defined before it is used.

So CFI is an outgrowth of the LLVM project work. It has been around for a few years and has been incorporated into various projects. Microsoft first incorporated it into Windows as Control Flow Guard (CFG) in Windows 8.1 Update 3 (KB3000850) in November 2014. Developers can add CFG to their programs by adding the /guard:cf linker flag before program linking in Visual Studio 2015 or later. As of Windows 10 Creators Update (Windows 10 version 1703), the Windows kernel is compiled with CFG.

The central idea is to catch and prevent malicious code reuse by checking the jump destination address of any data-driven indirect jumps. In an indirect jump, for example, existing code is jumping to another location which is specified by the data contained in a register. Since bad guys might be able to change a register's contents, they would then be able to execute an existing indirect jump to a location of their choosing.

Windows prevents this by maintaining a bitmap of allowable valid jump destinations and when a CFG application is being built, a bit of extra bitmap-checking code is added before each indirect jump instruction to verify that the location it is jumping to is a valid destination.

The Chromium web browser uses similar CFI technology under Linux where the addition of CFI has been benchmarked with an overhead of less than 1%. In that case the scheme had not been optimized for size, so it did inflate Chromium's size by 15%... but in today's “plenty of memory” computation world, the protection provided to a large attack surface such as a web browser probably merits this solution.

So our news is that Google's Pixel 3 is the first Android device to ship with CFI protecting its operating system kernel. Android 9.0 (Pie), released at the end of August, was the first Android OS version to feature some CFI support.

On June 28th earlier this year, Google wrote in the Android Developers blog:
https://android-developers.googleblog.com/2018/06/compiler-based-security-mitigations-in.html
"Compiler-based security mitigations in Android P"

Android's switch to LLVM/Clang as the default platform compiler in Android 7.0 opened up more possibilities for improving our defense-in-depth security posture. In the past couple of releases, we've rolled out additional compiler-based mitigations to make bugs harder to exploit and prevent certain types of bugs from becoming vulnerabilities. In Android P, we're expanding our existing compiler mitigations, which instrument runtime operations to fail safely when undefined behavior occurs. This post describes the new build system support for Control Flow Integrity and Integer Overflow Sanitization.

So, our takeaway here is: "It's in there" -- and it's good.

US GAO says: DOD Just Beginning to Grapple with Scale of Vulnerabilities
(And that's being kind.)
https://www.gao.gov/assets/700/694913.pdf

The GAO is the "US Government Accountability Office." We'll avoid, here, dwelling on the fact that the phrase "government accountability" is, itself, an oxymoron.

Why GAO Did This Study
DOD plans to spend about $1.66 trillion to develop its current portfolio of major weapon systems. Potential adversaries have developed advanced cyber-espionage and cyber-attack capabilities that target DOD systems.

What GAO Found
The Department of Defense (DOD) faces mounting challenges in protecting its weapon systems from increasingly sophisticated cyber threats. This state is due to the computerized nature of weapon systems; DOD's late start in prioritizing weapon systems cybersecurity; and DOD's nascent understanding of how to develop more secure weapon systems. DOD weapon systems are more software dependent and more networked than ever before.

Automation and connectivity are fundamental enablers of DOD's modern military capabilities. However, they make weapon systems more vulnerable to cyber attacks. Although GAO and others have warned of cyber risks for decades, until recently, DOD did not prioritize weapon systems cybersecurity. Finally, DOD is still determining how best to address weapon systems cybersecurity.

In operational testing, DOD routinely found mission-critical cyber vulnerabilities in systems that were under development, yet program officials GAO met with believed their systems were secure and discounted some test results as unrealistic.

Using relatively simple tools and techniques, testers were able to take control of systems and largely operate undetected, due in part to basic issues such as poor password management and unencrypted communications.

In addition, vulnerabilities that DOD is aware of likely represent a fraction of total vulnerabilities due to testing limitations. For example, not all programs have been tested and tests do not reflect the full range of threats.

So begins a 50-page report detailing nearly total disregard for cyber-security within the US Department of Defense.

Test Teams Easily Took Control

Test teams were able to defeat weapon systems cybersecurity controls meant to keep adversaries from gaining unauthorized access to the systems. In one case, it took a two-person test team just one hour to gain initial access to a weapon system and one day to gain full control of the system they were testing.

Some programs fared better than others. For example, one assessment found that the weapon system satisfactorily prevented unauthorized access by remote users, but not insiders and near-siders. Once they gained initial access, test teams were often able to move throughout a system, escalating their privileges until they had taken full or partial control of a system.

In one case, the test team took control of the operators' terminals. They could see, in real-time, what the operators were seeing on their screens and could manipulate the system. They were able to disrupt the system and observe how the operators responded. Another test team reported that they caused a pop-up message to appear on users' terminals instructing them to insert two quarters to continue operating.

The test reports indicated that test teams used simple tools and techniques to disrupt or access and take control of weapon systems. For example, in some cases, simply scanning a system caused parts of the system to shut down. One test had to be stopped due to safety concerns after the test team scanned the system.

Yes... in other words, as horribly insecure as our blue-box SoHo routers are, they are arguably more robust than the major United States Department of Defense weapons systems. Our tax dollars hard at work.

Program offices were aware of some of the weapon system vulnerabilities that test teams exploited because they had been identified in previous cybersecurity assessments. For example, one test report indicated that only 1 of 20 cyber vulnerabilities identified in a previous assessment had been corrected. The test team exploited the same vulnerabilities to gain control of the system. When asked why vulnerabilities had not been addressed, program officials said they had identified a solution, but for some reason it had not been implemented. They attributed it to contractor error. Another test report indicated that the test team exploited 10 vulnerabilities that had been identified in previous assessments.

The EFF versus the widespread use of facial recognition
https://www.eff.org/deeplinks/2018/10/chicago-should-reject-proposal-private-sector-face-surveillance

A proposed amendment to the Chicago municipal code would allow businesses to use face surveillance systems that could invade biometric and location privacy, and violate a pioneering state privacy law adopted by Illinois a decade ago. EFF joined a letter with several allied privacy organizations explaining our concerns, which include issues with both the proposed law and the invasive technology it would irresponsibly expand.

At its core, facial recognition technology is an extraordinary menace to our digital liberties. Unchecked, the expanding proliferation of surveillance cameras, coupled with constant improvements in facial recognition technology, can create a surveillance infrastructure that the government and big companies can use to track everywhere we go in public places, including who we are with and what we are doing.

This system will deter law-abiding people from exercising their First Amendment rights in public places. Given continued inaccuracies in facial recognition systems, many people will be falsely identified as dangerous or wanted on warrants, which will subject them to unwanted—and often dangerous—interactions with law enforcement. This system will disparately burden people of color, who suffer a higher "false positive" rate due to additional flaws in these emerging systems.

In short, police should not be using facial recognition technology at all. Nor should businesses that wire their surveillance cameras into police spying networks.

Moreover, the Chicago ordinance would violate the Illinois Biometric Information Privacy Act (BIPA). This state law, adopted by Illinois statewide in 2008, is a groundbreaking measure that set a national standard. It requires companies to gain informed, opt-in consent from any individual before collecting biometric information from that person, or disclosing it to a third party. It also requires companies to store biometric information securely, sets a three-year limit on retaining information before it must be deleted, and empowers individuals whose rights are violated to enforce its provisions in court.

Having overcome several previous attempts to rescind or water down its requirements at the state level, BIPA now faces a new threat in a recently proposed municipal amendment in Chicago. The proposal to add a section on "Face Geometry Data" to the city's municipal code would allow businesses to use controversial and discriminatory face surveillance systems pursuant to licensing agreements with the Chicago Police Department.

---- It goes on, but we get the idea.

Once again we have a situation where technological advancement and progress is creating new capabilities that our established legal frameworks and assumptions have not been updated to address. When Apple's iPhone X introduced FaceID a great deal of concern was raised about its privacy implications and Apple went out of its way to explain how it was purely a local recognition and unlock capability. I'm very glad we have the EFF watching our backs (and not our faces.)

DuckDuckGo Is Now Receiving More Than 30 Million Searches in a Single Day
https://www.bleepingcomputer.com/news/technology/duckduckgo-is-now-receiving-more-than-30-million-searches-in-a-single-day/

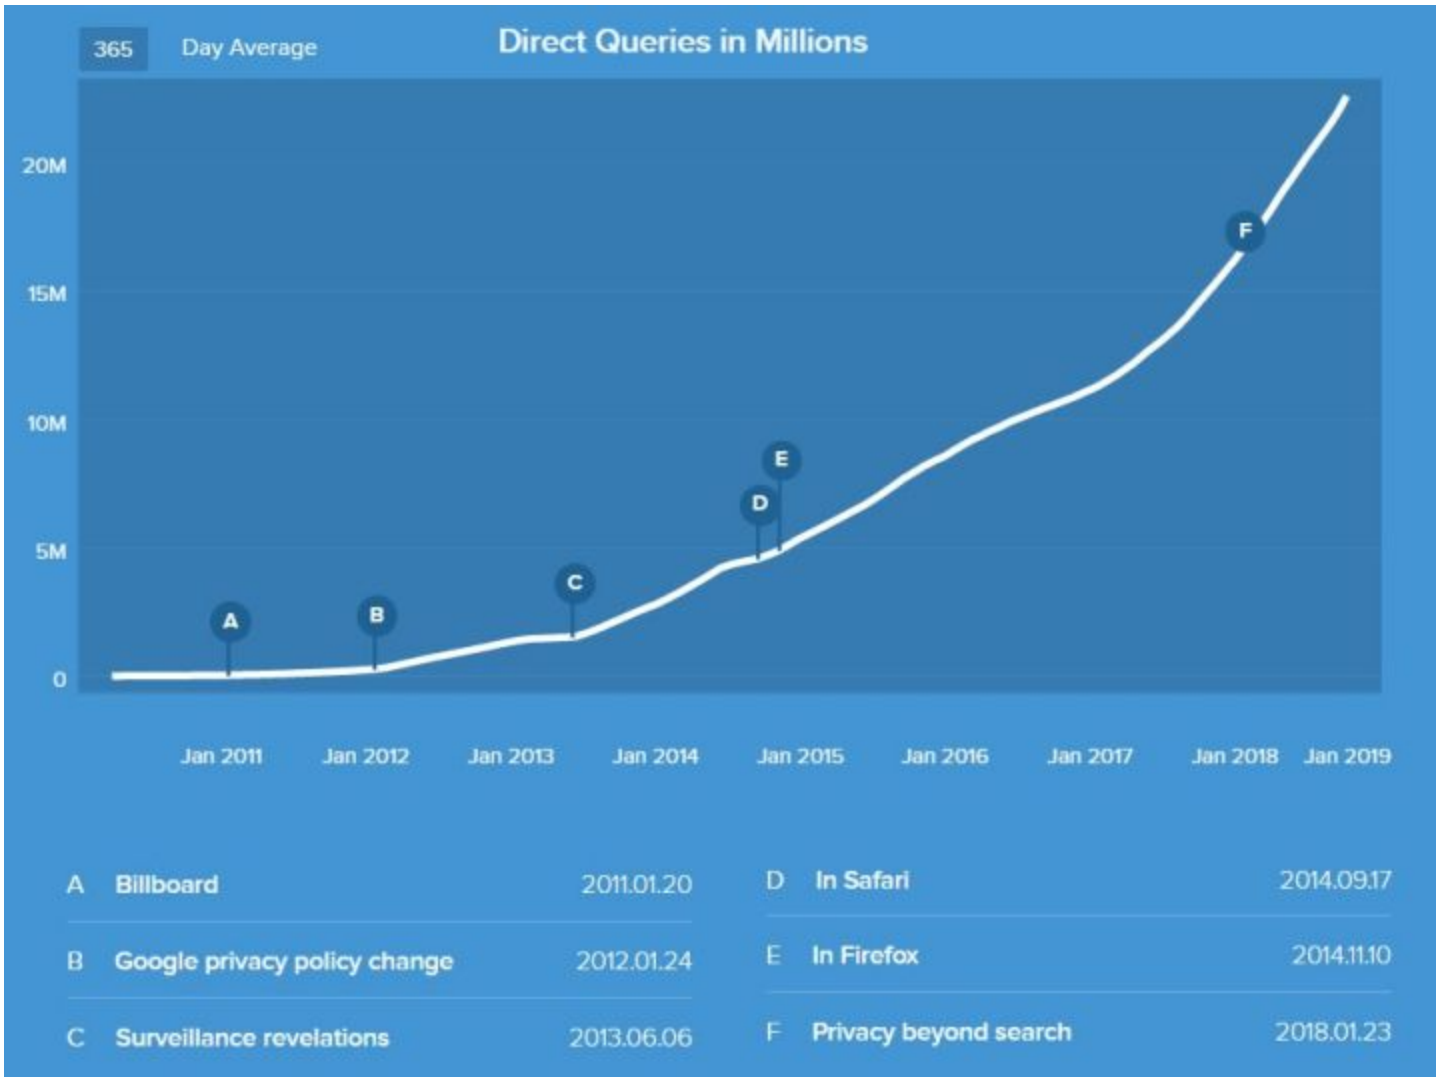

https://duckduckgo.com/traffic "The search engine that doesn't track you."

Google links are all redirection through Google to allow them to monitor what every Google search user clicks on. DuckDuckGo links are all direct links to the destination.

https://donttrack.us/ ← Wonderful! (See the 2nd page!)

PHP 5.x branch support ending in 10 weeks, at the end of this (2018) year…

According to statistics from W3Techs, today roughly 78.9% of all Internet sites run on PHP.

On December 31, 2018, security support for PHP 5.6.x will officially cease, marking the end of all support for any version of the ancient PHP 5.x branch.

This means that starting with next year, around 62 percent of all Internet sites still running a PHP 5.x version will stop receiving security updates for their server and website's underlying technology, exposing hundreds of millions of websites, if not more, to serious security risks. If a hacker finds a vulnerability in PHP after the New Year, lots of sites and users would be at risk.

MANY PHP-based systems such a Wordpress and Joomla set their minimum requirements at PHP v5.2 and v5.3, respectively. This allows website administrators to remain where they are.

Real-World Protection Test September 2018 – Factsheet
https://www.av-comparatives.org/tests/real-world-protection-test-september-2018-factsheet/

Our Real-World Protection Test is currently the most comprehensive and complex test available, using a large number of test cases. Currently, we are running this test under Microsoft Windows 10 Pro RS4 64 Bit with up-to-date third-party software (such as Adobe Flash, Adobe Acrobat Reader, Java, etc.). Due to this, finding in-the-field working exploits and running malware is much more challenging than e.g. under an non-up-to-date system with unpatched/vulnerable third-party applications.

The results are based on the test set of **197** live test cases (malicious URLs found in the field), consisting of working exploits (i.e. drive-by downloads) and URLs pointing directly to malware. Thus exactly the same infection vectors are used as a typical user would experience in everyday life. The test-cases used cover a wide range of current malicious sites and provide insights into the protection given by the various products (using **all** their protection features) while surfing the web.

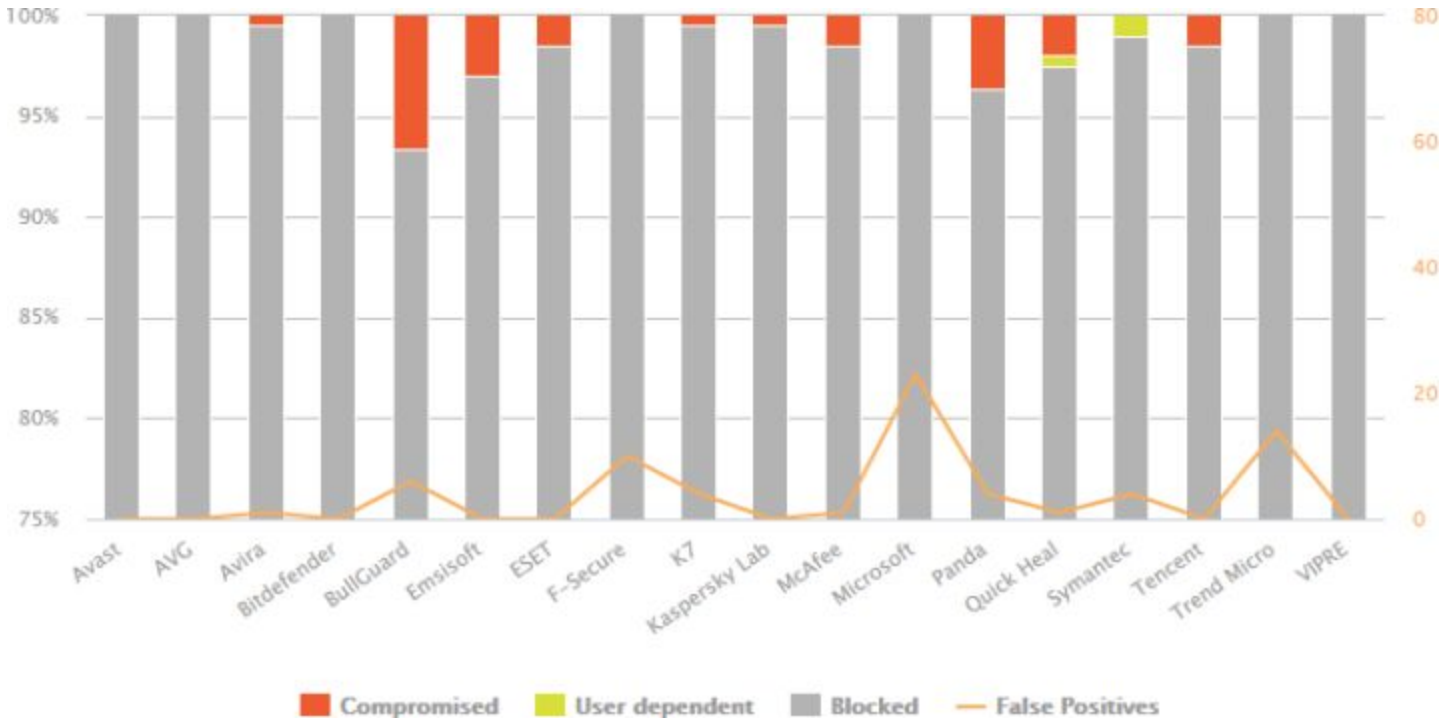

Closing The Loop

"ÞeJC™" @cronhan (Johan (pronounced "Yo-han"))
Hi! Listener of your show for about a year now, maybe longer. What are your thoughts on OCSP? Had an issue the other day where I couldn't access my ISP website via FF or Edge; could in Chrome. Turned off OCSP validation in FF and it worked. Keep the setting off or turn back on? Thx

Timo Grün @Khoji
On Bloomberg and SuperMicro: Why put the chips ON the boards where they can be found? Modern motherboards have so many layers and the chips are so small you could easily sandwich them INTO the board. Good luck finding that.

Sultan of Saki @EpcotJosh
Hi Steve. A friend recently asked me how they should go about selling some used hard drives on eBay. I explained to him the importance of wiping them, and as I was doing so I started thinking about how drives can "swap out" bad sectors for spares. It occurred to me that it is possible that if the drive was ever used unencrypted that any data in sectors that had been "swapped out" would thereafter be completely inaccessible to the OS or application layers, thus making it impossible to wipe the data they contain. If my analysis of this issue is correct, then this is yet another reason that users should always ensure that they encrypt their drives immediately upon installation and in no case should they copy data to the drive until the encryption is complete. Do you think my thinking is reasonable here or am I missing something? Thanks! Josh Fenton

Gary Napier @_GaryNapier
Hey Steve, thanks for all the great info! Do you know of any way to check your router to see if you are infected with VPNFilter? Thanks!

Mike Estes @mle_ii
Forgot to message you before and after last week's episode. You'll be happy to hear that the NetGear router I am using that was in the list of being vulnerable to VPN Filter not only got another update, it also got an update to start doing automatic updates which I opted into. Seems they're getting it, well at least getting part of it. Thank you and keep up the great work. Love the show and listening to you talk on security.

John McAfee @officialmcafee
The "Presidential alerts": they are capable of accessing the E911 chip in your phones - giving them full access to your location, microphone, camera and every function of your phone. This not a rant, this is from me, still one of the leading cybersecurity experts. Wake up people!

Good Samaritans

Exactly how severe is the problem?
http://www.theamericanconsumer.org/wp-content/uploads/2018/09/FINAL-Wi-Fi-Router-Vulnerabilities.pdf

A new study conducted by the American Consumer Institute (ACI), a US consumer nonprofit, has found that five out of six home routers are inadequately updated for security flaws, leaving the devices, and indirectly their users, vulnerable to hacking.

The study (see PDF above) analyzed a sample of 186 SOHO (small office/home office) Wi-Fi routers from 14 different vendors with a presence on the US market. ACI looked at the firmware version the routers were running and searched public vulnerabilities databases for known security flaws affecting each device's firmware.

"In total, there were a staggering 32,003 known vulnerabilities found in the sample. Our analysis shows that of the 186 sampled routers, 155 (83 percent) were found to have vulnerabilities to potential cyberattacks, in the router firmware, with an average of 172 vulnerabilities per router, or 186 vulnerabilities per router for the identified 155 routers," ACI experts said.

Of the total 32,003 security flaws, more than a quarter were vulnerabilities that received the two highest severity ratings of "critical" and "high-risk" respectively.

"Our analysis shows that, on average, routers contained 12 critical vulnerabilities and 36 high-risk vulnerabilities, across the entire sample," researchers said.

These are staggeringly large numbers.

So… **Oct 12th: A mysterious vigilante grey-hat hacker is patching people's outdated MikroTik routers**

https://www.zdnet.com/article/a-mysterious-grey-hat-is-patching-peoples-outdated-mikrotik-routers/

First of all: We stipulate that the BEST solution would be:

1. No bugs in widely popular router firmware.
2. Reduced attack surface by default disabling remote services.
   (Require users to deliberately enable remote access features rather than requiring them to deliberately disable remote access features.)

3. Default to:
    1. Notify user of available firmware updates and obtain permission to pdate firmware.
       That way, if something goes sideways a local human is in the loop to fix.
    2. Allow user to enable pushing auto-updates with a firmware rollback on failure.

Adding firewall rules to close wayward vulnerable ports

It would be better to update the firmware to remove vulnerabilities

Reports of infuriated users

Without permission, the user may have been using those open ports

Notification without modification?