Security Now! #622 - 08-01-17

Hack the Vote

This week on Security Now!

This week we look at the expected DefCon fallout including the hacking of US election voting machines, Microsoft's enhanced bug bounty program, the wormification of the Broadcom WiFi firmware flaw, the worries when autonomous AI agents begin speaking in their own language which we cannot understand, Apple's pulling VPN clients from its Chinese app store, a follow-up on iRobot's floorplan mapping intentions, some new on the Chrome browser front, the 18th Vault-7 Wikileaks dump, and some closing-the-loop feedback from our terrific podcast followers.

"Ineffective Countermeasures"

Security News

Hack The Vote
During the recently completed DefCon cybersecurity conference last week, several hackers managed to hack into multiple US voting machines... in some cases within minutes, and in others within a few hours.

For the first time, but likely not the last, DefCon hosted a "Voting Machine Village" during which the conferences tech-savvy attendees tried -- and succeeded -- to hack many commercial voting machine systems and help catch vulnerabilities.

This year's DefCon Voting Machine Village provided 30 different pieces of voting equipment used in American elections, including a Sequoia AVC Edge, ES&S iVotronic, AccuVote TSX, WinVote, and Diebold Expresspoll 4000 voting machines.

The executive summary is that the conference acquired 30 machines for hackers to poke at... and every voting machine in the village was hacked to varying degrees, often requiring less than 90 minutes to compromise a machine.

The DefCon hackers took complete control of an e-poll book, which is currently in use in dozens of states where voters sign in and receive their ballots. They also discovered and exploited significant security flaws in the AccuVote TSX, currently in use in 19 states, and the Sequoia AVC Edge, used in 13 states. And another hacker broke into the hardware and firmware of the Diebold TSX voting machine.

Though somewhat less surprising, the WinVote voting machine had long been removed from use due its vulnerabilities, which were again confirmed... though it was once widely used. They found a remote access vulnerability in WinVote's OS which exposed real election data that was still stored in the machine.

And another hacker hacked into the Express-Pollbook system and exposed the internal data structure via a known OpenSSL vulnerability (CVE-2011-4109), allowing anyone to carry out remote attacks.

Jake Braun, a cybersecurity expert at the University of Chicago, who sold DEF CON's founder Jeff Moss on the idea of creating the Voting Machine Village said "Without question, our voting systems are weak and susceptible. Thanks to the contributors of the hacker community today, we have uncovered even more about exactly how."

Results are being assembled and tallied and will be assembled into a final follow-up report on Github:
https://github.com/josephlhall/dc25-votingvillage-report

Harri Hursti, cofounder of Nordic Innovation Labs, one of the event's organizers said: “The Village was announced at the last minute. But people were active in the forums, looking to understand the problem. The changes have to start somewhere. This year it’s in this room, next year it will be a bigger room.”

Eric Hodge, director of consulting at CyberScout, a consultant for Kentucky's Board of Elections said: "The best possible outcome is that the village results in a book of vulnerabilities to share with the FEC (Federal Election Commission), states, and other firms like ours"

DEF CON's Voting Machine Village was the first time most researchers had ever had access to voting machines. That's what has change. There is no way that any state or county government should be allowed to spend taxpayer money on machines which have not been independently audited by independent security researchers. Rather than treating their machines like proprietary closed boxes, voting machine manufacturers should gleefully turn their machines over to every independent security researcher they can find to harden their offerings security. Then purchasers should have candidate purchases again independently vetted by still other independent researchers.

Every lesson we have learned on this podcast through its nearly twelve years has informed us that systems much first be designed to be secure. But also that such design is insufficient for anything but a well-intentioned claim of security. The only way to have any practical assurance of delivered security is to have multiple, unaffiliated, researchers, armed with the same motivation and resources as determined attackers, attempt to breach a system's security.

I am again reminded of Steve Ballmer's pre-WinXP release, where he was screaming into a microphone on stage that XP was a totally secure operating system -- before its release. And as we know, XP was initially a security disaster. As I've often said, the security of a system must first be designed in... but any systems actual delivered security can only be proven afterwards.

At the moment we have commercial voting machine companies acting just like Steve Ballmer, touting their "military grade" encryption and everything else... while at the same time actively preventing any independent auditing of those claims. The ONLY WAY to ever obtain voting machine security will be for this paradigm to change so that independent fully open security auditing becomes part of the purchase cycle.

DEFCON reveals an SMB server takedown that Microsoft has said they won't be fixing.
Eight years ago, back in 2009, an attack tool known as "Slow Loris" operated over HTTP web connections. It was a server-side resource depletion attack which exhausted a server's incoming connection-handling capacity while requiring very little bandwidth from the attacker. Unlike traditional DoS attacks then and now, even one attacking machine could bring a remote website to its knees, preventing legitimate visitors from obtaining access.

Today, as revealed at last week's DEFCON, we have a similar server-side resource depletion attack using that troublesome version 1 of Microsoft's still-supported, yet long obsolete, file and printer sharing protocol known as SMB, for Server Message Blocks.

Sean Dillon of RiskSense was among the first researchers to analyze EternalBlue, the leaked NSA SMB exploit that was used to spread the WannaCry ransomware attack. It was during that analysis that Dillon uncovered this issue.

"While working on EternalBlue, we observed a pattern in the way memory allocations were done on the non-paged pool of the Windows kernel. The non-paged pool is memory that has to be

reserved in physical RAM; it can't be swapped out," Dillon explained. "That's the most precious pool of memory on the system. We figured out how to exhaust that pool, even on servers that are very beefy, even 128 GB of memory. We can take that down with a Raspberry Pi."

Like Slow Loris before it, this new attack, dubbed "SMBloris", leverages a memory handling bug that could be exploited by attackers to shut down big web servers with small computers.

However, while it was initially believed that attackers could only exploit the SMBloris vulnerability if the target machine has SMBv1 exposed to the Internet, but that now appears to be hopeful thinking. The Register later updated their coverage to say: "According to Microsoft's SMB supremo Ned Pyle, SMBLoris affects ALL VERSIONS of SMB – not v1 as first thought – because it all happens so early on in the connection." So this problem will NOT be going away even after the forthcoming Win10 fall creator's update disables support for SMBv1.

NBSS is the NetBIOS Session Service protocol. And every connection to it statically allocates 128 KB of page-locked non-swappable kernel memory which is freed only when the connection is closed. Connections will be proactively closed after 30 seconds of no activity.

But with more than 64 thousand TCP ports available, attackers can force the allocation of more than 8 GB. And since IPv4 and IPv6 can both be used, it's possible to tie up 16 GB. And if a second source IP were available, the memory commit could be doubled to 32 GB of RAM using just two IPs. Once the attack has triggered memory saturation the server must be rebooted to restore normal operation.

Sean Dillon of RiskSense, the discovered of the attack noted that: "A Raspberry Pi could take down the beefiest server." (Using 20 lines of Python code.)

We know that many systems have SMB exposed to the Internet. But well-configured public-facing Internet servers should never have SMB exposed, most ISPs still block and drop any traffic attempting to come in over those ports to their subscribers, any NAT router will block incoming traffic, and OS firewalls will do so too.

Sean said: "I think Microsoft's problem is that it would not be easy to fix; it's the way they've done SMB memory allocation for over 20 years. So everything relies on the fact the client says 'I have a buffer that I'm sending that's this big.' So then the server reserves that much memory so it can handle the anticipated incoming data. What we did with this attack was to simply say 'I have a huge buffer' and never send the buffer. There are many components of the protocol support that rely on the fact that buffer is already allocated and the size is already known."

So... we have an old bug that's present in every version of Windows. This means that attackers will likely attack those machines that do have SMB enabled and exposed to the public Internet. And while this COULD in theory be leveraged into a more powerful attack to force a reboot where one is needed, this is unlikely to affect our properly configured podcast listeners.

Meanwhile... Microsoft has announced a new Windows Bug Bounty Program
Last Thursday, Microsoft announced a new Windows Bounty Program that will pay researchers up to $250,000 for finding and disclosing security vulnerabilities.

The bug bounty program will focuse upon a few key areas: Hyper-V, Mitigation bypass, Windows Defender Application Guard, Microsoft Edge, and all features made available via the Windows Insider Program. Payouts depend on where a vulnerability is found and how severe it is and can range from $500 for a vulnerability in Edge all the way up to a quarter of a million dollars for a critical vulnerability in Hyper-V.

In their announcement Microsoft said: "Any critical or important class remote code execution, elevation of privilege, or design flaws that compromises a customer’s privacy and security will receive a bounty."

And, interestingly, this also applies to previously-reported bugs: Microsoft said it will pay "a maximum of 10% of the highest amount" discoverers would have received if the discovery was fresh. Thus Microsoft is working to encourage researchers to disclose everything instead of sitting on vulnerabilities because they don't know if they're new.

Microsoft lays out the various categories and payout ranges on this page:
https://technet.microsoft.com/en-us/security/dn425036

DefCon's "BroadPwn" wormifies the Broadcom WiFi firmware bug
Until both Google and Apple both issued patches last month, an estimated 1 billion devices were vulnerable to this new, remotely exploitable worm attack which Exodus Intelligence's Nitay Artenstein has dubbed "Broadpwn."

And this is the way it always goes: First a device is crashed, then the vulnerability is carefully examined and weaponized into something far more functional and useful to attackers.

At the Black Hat security conference, Nitay demonstrated proof-of-concept attack code that exploited this vulnerability we've been discussing through the past month which affected the widely and pervasively used BCM43xx family of WiFi chips manufactured by Broadcom.

Nitay's attack fills the airwaves surrounding any compromised device with probes requesting low-level (no user alert or action required) connection requests to any and all nearby mobile smartphones. When the specially devised requests reach a device using the BCM43xx family of Wi-Fi chipsets, the attack rewrites the firmware that controls the chip. The compromised chip then sends the same malicious packets to other vulnerable devices, setting off a potential chain reaction.

Nitay wrote: "Broadpwn is a fully remote attack against Broadcom’s BCM43xx family of Wi-Fi chipsets, which allows for code execution on the main application processor in both Android and iOS. It is based on an unusually powerful 0-day that allowed us to leverage it into a reliable, fully remote exploit." Nitay said his attack worked on a wide range of phones, including all iPhones since the iPhone 5, Google's Nexus 5, 6, 6X and 6P models, Samsung Note 3 devices, and Samsung Galaxy devices from S3 to S8.

This is another perfect substantiation of what we've been saying on this podcast now for several years: it is utterly unsafe to be using any iOS- or Android-based device that is not receiving regular security updates. And note that Android is not alone here, since Apple also eventually abandons their older, yet still functional, iOS devices once their hardware can no longer run the latest version of iOS.

Nitay Artenstein's fascinating blog posting at Exodus Intelligence contains every detail:
https://blog.exodusintel.com/2017/07/26/broadpwn/

Facebook AI Agents negotiating in their own language...
FoxNews: Facebook engineers panic, pull plug on AI after bots develop their own language
BGR: Facebook AIs develop own language and are immediately shutdown.

Many news outlets covered this story and picked up on each other's coverage. I tracked down the original coverage which was much drier and far more accurate:
https://www.fastcodesign.com/90132632/ai-is-inventing-its-own-perfect-languages-should-we-let-it

AI Is Inventing Languages Humans Can't Understand. Should We Stop It?
Researchers at Facebook realized their bots were chattering in a new language.
Then they stopped it.

Bob: "I can can I I everything else."

Alice: "Balls have zero to me to me to me to me to me to me to me to me to."

That passage looks like nonsense, but this "nonsense" was the discussion of what might be the most sophisticated negotiation software on the planet? Negotiation software that had learned, and evolved, to get the best deal possible with more speed and efficiency–and perhaps, hidden nuance–than we're able to perceive.

This conversation occurred between two AI agents developed inside Facebook. At first, they were speaking to each other in plain old English. But then researchers realized they'd made a mistake in programming.

Dhruv Batra, at Facebook AI Research (FAIR) as a visiting research scientist from Georgia Tech said: "There was no reward for sticking with the English language as the AIs conversed." The two AI agents were competing to get the best deal, an effective strategy for sharpening the operation of AI by pitting them against each other in what is known as a "generative adversarial network." In this case, neither was offered any incentive for speaking as a normal person would. So as they grew, they began to diverge from English, eventually rearranging legible words into seemingly nonsensical sentences. But sentences they each understood.

Batra, speaking to a now-predictable phenomenon that's been observed again, and again, and again, said: "AI agents will drift away from understandable language, inventing more efficient codewords for themselves.

So... should we let our software do the same thing? Should we allow AI to evolve its dialects for specific tasks that involve speaking to other AIs? To essentially gossip out of our earshot? Maybe; it offers us the possibility of a more interoperable world, a more perfect place where iPhones talk to refrigerators that talk to your car without a second thought. The tradeoff is that we, as humanity, would have no clue what those machines were actually saying to one another.

Mike Lewis, research scientist at FAIR said that Facebook ultimately opted to require its negotiation bots to speak in plain old English. "Our interest was having bots who could talk to people." And Facebook isn't alone in that perspective. Microsoft has also indicated that it's more focused upon human-to-computer speech. Meanwhile, Google, Amazon, and Apple are all also focusing incredible energies on developing conversational personalities for human consumption. They're the next wave of user interface, like the mouse and keyboard for the AI era.

The other issue, as Facebook admits, is that it has no way of truly understanding any divergent computer language. Batra says: "It's important to remember, there aren't bilingual speakers of AI and human languages. We already don't generally understand how complex AIs think because we can't see inside their thought process. Adding AI-to-AI conversations to this scenario would only make that problem worse.

Facebook's discussion of AI negotiation.
https://code.facebook.com/posts/1686672014972296/deal-or-no-deal-training-ai-bots-to-negotiate/

Over the weekend, Apple pulled many iOS VPN clients from its Chinese App Store.
The makers of those VPN applications were up in arms over Apple's "capitulation" to pressure from the Chinese government. They were claiming that this was a human rights issue and were "disappointed" in Apple. But I think that's a load of nonsense.

Greater China is Apple's largest market outside the United States. And Apple is a super-successful publicly-held commercial for-profit company whose device application model is that of a closed and customer-protecting application ecosystem.

And, yes... that means there's a tradeoff. If you want an open ecosystem there are plenty of them freely available. Get a PC or jailbreak an Android phone and sideload whatever you want... and take your chances.

Apple knows quite well that if they are going to operate in China they can only do so with the permission and approval of the Chinese government. And if I were in China I would definitely want the significantly increased security of using an Apple iOS device rather than any other non-Apple solution.

Yes, it's closed... but it's also closed, as much as is possible, to malware... and for the moment we also believe that it's likely closed to eavesdropping.

iRobot walks back their plans to sell the users' homes' floor plans.
Following a firestorm of customer outrage over the news, iRobot's CEO Colin Angle now says that his comments to Reuters were misinterpreted and that iRobot has no plans and would never make such data available to third parties.

Wondering what was "misinterpreted, I went back to our coverage last week and found Colin's quote from Reuters: "There's an entire ecosystem of things and services that the smart home can deliver once you have a rich map of the home that the user has allowed to be shared."

Google's Chrome browser to begin tightening the screws on embedded iFrames
iFrame is short for "Inline Frame" - essentially a web page inside a web page... and this is a VERY POWERFUL capability since it allows one website's page to transparently host any other within itself.

Consequently, and not surprisingly, iFrames have historically been an outsized vector for malware and browser exploitation.

Therefore, beginning with Chrome 63, which is expected near the end of this (2017) year, iFrames will be restricted in what they may contain, with restricted permissions needing to be explicitly allowed:

An "Allow" term will be added that can contain any of the following values:

- geolocation
- microphone
- camera
- speakers
- midi
- encrypted-media

Those features will be denied by default, and the embedding page will need to explicitly enable each feature subset which it wishes to allow the contained content to access.

In more good new on the Chrome front:
Users running the newest version Chrome (60) now have very easy access to more details about security certificates, and their validity, by clicking on the padlock in the left side of the address bar.

When explicitly enabled, a new "Certificate" item becomes the first item on the drop-down list. A "Valid" link, when clicked, displays the page's certificate, making its access much quicker and easier. For the time being, this added certificate UI is opt-in.

Go to 'chrome://flags/#show-cert-link' to enable the 'Show certificate link' functionality in the browser.

MANY flags. It's very near the bottom. Enable the item and click the "relaunch" button.

WikiLeaks 18th Vault-7 dump:

Not to be forgotten, in their 18th Vault-7 document dump, last week WikiLeaks appeared to reveal details of three more CIA hacking and implant tools named Achilles, SeaPea and Aeris targeted at Mac and Linux systems.

- *"Achilles"* is a tool to subvert Mac OS X Disk Images. It allows CIA operators to combine malicious Trojan applications with a legitimate Mac OS app into a disk image installer (.DMG) file. The tool which binds the pieces together is written in BASH shell script, giving operators the opportunity to run the appended malware as desired. When the unsuspecting targeted user downloads an infected disk image on their Apple computer, opens and installs the software, the malicious executables would also run in the background.

  Significantly, Achilles also detaches itself so that afterwards, all the traces of the Achilles tool are "removed securely" from the downloaded application so that the file would "exactly resemble" the original legitimate app, un-trojaned application, making it hard for the investigators and antivirus software to detect the initial infection vector.

  The primary interest in this is not any particular "high techness" but, assuming that these documents are legitimate and were leaked from the CIA as claimed, it's at least informative that this sort of targeted attacking is going on.

- *"SeaPea"* is a stealth rootkit for Mac OS X systems. As such, like any rootkit, SeaPea would provide CIA operators with a stealth tool with launching capabilities that hides important files, processes and socket connections from the users, allowing them to access Macs without the targeted victim's knowledge. SeaPea requires transient root access to be installed on a target Mac computer and cannot be removed unless the startup disk is reformatted or the infected Mac is upgraded to the next version of the operating system.

- *"Aeris"* is an automated implant targeted at Linux machines. Written in C, it is designed to open backdoors in Linux systems including those from Debian, CentOS, Red Hat, along with the FreeBSD and Solaris UNIX operating systems. Aeris provides build scripts which allow CIA operators to generate customized effect depending upon their needs. It supports automated file exfiltration, configurable beacon interval and jitter, HTTPS and SMTP protocol support — all with TLS encrypted communications with mutual authentication. It provides fully secure end-to-end encrypted exfiltration with a structured command and control that's similar to that used by several Windows implants.

These are all workman-like tools. They don't leverage 0-day or known exploits. They simply rely upon somehow arranging to get the malicious code to run once on a target's machine.

SpinRite

Blog Posting: Universal fix for windows KSOD

Ever had your Windows installation inexplicably die leaving your computer unusable without a fix? I have – more times than I'd like to count. The last time this happened was yesterday when Windows 7 would only boot into a black screen with a movable cursor, also known as the blacK Screen Of Death. It was a serious case considering none of the safe modes or repair function in the Windows boot options would work; each option would universally end in either a KSOD or the classic BSOD after hanging on aswRvrt.sys during safeboot. After exhaustively eliminating all possible "regular" fixes that were available on the internet, I decided it was time for the big guns: Steve Gibson's Spinrite.

Prior to trying Spinrite I first tried Kaspersky's Rescue Disc 10 which was entirely useless for my case. After booting from the rescue USB dongle I would always get a "Missing Operating System" error in the boot screen. Not reassuring. I have long been a fan of Steve Gibson's Security Now podcast, which is why I knew of the tool. I knew the tool would be one of the few things that might do the trick, so I gave it a shot. After about an hour running Spinrite 6, and a few reboots later, my Windows 7 installation was working perfectly as if nothing had ever happened. Spinrite saved the day.

TL;DR Spinrite saved my machine from a perpetual and otherwise unbeatable KSOD scenario and my guess is that if you are having KSOD problems then Spinrite is one of few things that might help you too.

Closing The Loop

Logan Rogers (@techieg33k)
@SGgrc Are you just using an #S3 Bucket from @awscloud for your backup that you talk about on @SecurityNow ? Looks like a cheap cloud backup.

ugly b0b (@uglyb0b)
@SGgrc I'm curious as to what software you use to backup to aws?

"S3" Standalone Windows .EXE command line utility for Amazon S3 & EC2
http://s3.codeplex.com/

A Windows command-line utility for Amazon's S3 & EC2 web services that requires no installation, is a single .EXE file with no DLLs, and requires only .NET 2.0 or Mono, so will work on a plain vanilla Windows 2003 installation.

Key Features:

- Efficiently uploads and downloads large numbers of files (or whole directories) between Amazon S3 and Windows PCs.
- Everything is in one .EXE. Nothing to install or configure, just download it where it's needed and run.

- Doesn't require anything except .NET 2.0 or Mono [which version?] which you already have on all your machines (don't you?).
- Works well in an automated backup solution or as an ad-hoc system administration tool.
- Can split large files into chunks for upload without creating any temporary files on disk.
- Fast parallel file transfers (coming soon).
- Can use HTTP HEAD command to quickly determine which files don't need to be uploaded because they haven't been updated (/sync).
- Support for EC2 operations as well.
- Free & open source. There is no paid version.

Steve Whisenhant (@stevewhisenhant)
@SGgrc, I drive Uber & I'm considering enabling hot spot on my phone as perk for riders. What are the risks to me assuming WPA2 & PW posted in the car?

The Best Mobile Hotspots of 2017
https://www.pcmag.com/article2/0,2817,2400503,00.asp

Trust No-One (@twust)
@SGgrc @LastPass 2FA is only as strong as a cleartext link emailed to u when u click "I lost my 2FA"

The important lesson here is "the weakest link" principle: No security system is stronger than its weakest link. So no authentication system can be stronger than the strength of what's required to bypass it.

This is the tradeoff that SQRL makes...

Rex Moncrief (@smartergeek)
Hey @SGgrc any feedback on how this “retrieves” data from threema?
http://www.zdnet.com/article/this-android-spyware-can-record-calls-take-screenshots-and-video-targets-gmail-linkedin-snapchat/

Chris_shaw (@St0ble)
@SGgrc does printing a PDF to a PDF strip malcious content from the file?

~30~