Security Now! #563 - 06-07-16

IoT Infancy (pt.2)

This week on Security Now!

- A "Reality-Check" timeout.
- A new 0-day Windows exploit on the market
- A truly horrifying (and clever) chip-level exploit
- Yesterday's monthly Android Security Update
- A sad side-effect of the GWX push
- The LinkedIn breach apparently bites Mark Zuckerberg
- Facebook plans to offer optional encryption for Messenger
- Five things that give self-driving cars headaches
- A follow-up on SQRL's authentication management
- Some miscellany... and...
- Some truly horrifying details of Internet-connected Baby Monitor Implementations

Reality Check Timeout

The crucial difference between security FEATURES and actual security.

- "Windows 10 has added security features" is absolutely true.
- "Windows 10 is more secure" is utter nonsense.
- Steve Ballmer prancing around on stage declaring, before it's release, that Windows XP was the most secure operating system ever. It turned out to be the LEAST secure OS ever. Why? Because of all the new stuff they added.
- When you add more code, you add more vulnerabilities.
- Leo was talking about "Debian's Stable Track"...
- Don't even mess with the screen saver.

Windows 10 is NOT a "new" operating system.

- Microsoft doesn't create new operating systems.
  - They make new **LOOKING** operating systems, merely to sell new stuff.
- We KNOW it's not "new" in two ways:
  - **One:** If you drill down into the UI a few levels you encounter dialog boxes that haven't changed in 20 years. That's the same old code that's been left alone because it's utilitarian and no sexy.
  - **Two:** Because every time a flaw is found, it needs to be patched all the way back in time, across all previous versions of the OS, to the beginning -- why?? Because the same code is in use now as was in use then. If it was actually a new operating system there would be a complete disconnect among security flaws. Instead... every flaw exists in EVERY previous Windows OS version.

Security News

Most Windows versions suffer from new zero-day exploit
http://www.winbeta.org/news/windows-versions-suffer-new-zero-day-exploit

- A seller by the name of "BuggiCorp" has been offering a pervasive Windows zero-day exploit for $95,000. It's a Local Privilege Escalation (LPE).

- English translation from the Russian posting:
  Exploit for local privilege escalation (LPE) for a 0day vulnerability in win32k.sys. The vulnerability exists in the incorrect handling of window objects, which have certain properties, and [the vulnerability] exists in all OS [versions], starting from Windows 2000. [The] exploit is implemented for all OS architectures (x86 and x64), starting from Windows XP, including Windows Server versions, and up to current variants of Windows 10. The vulnerability is of "write-what-where" type, and as such allows one to write a certain value to any address [in memory], which is sufficient for a full exploit. The exploit successfully escapes from application containment, bypassing (more precisely: doesn't get affected at all [by]) all existing protection mechanisms such as ASLR, DEP, SMEP, etc. [The exploit] relies solely on the KERNEL32 and USER32 libraries [DLLs]. The [source code] project of the exploit and a demo example are written in C and assembly with MSVC 2005. The output is a "lib"-file which can later be linked to any other code, and [additional output from the source code project] is a demo EXE file which launches CMD EXE and escalates the privileges to SYSTEM account. The resulting EXE file size is between 7KB to 12KB depending on OS architecture. The exploit was tested on all versions of Windows, starting from XP, and on at least 20 different variants of Windows OS, including Windows Server versions.

- Trustwave, who found and reported the exploit offer explained the value of a Local Privilege Escalation nicely, writing: "Although such an exploit can't provide the initial infection vector like a Remote Code Execution (RCE) would, it is still a very much needed puzzle piece in the overall infection process. For instance, an LPE exploit paired with a client-side RCE exploit can allow an attacker to escape an application that implements sandbox protection (For example Google Chrome, Adobe Reader, etc…).

  Moreover, an LPE exploit provides the means to persist on an infected machine, which is a crucial aspect when considering APTs (Advanced Persistent Threats). In general terms, this exploit can be leveraged in almost any kind of attack scenario.

- The sales offering includes two videos showing the exploit in use. The seller wants to make it very clear that this is a real vulnerability, not a scam.

- Putting this into context: It is quite rare to find a 0-day exploit of any kind offered for sale. Trustwave wrote: "Finding a zero day listed in between these fairly common offerings is definitely an anomaly. It goes to show that zero days are coming out of the shadows and are fast becoming a commodity for the masses, a worrying trend indeed."

This ‘Demonically Clever’ Backdoor Hides In a Tiny Slice of a Computer Chip
https://www.wired.com/2016/06/demonically-clever-backdoor-hides-inside-computer-chip

- Andy Greenberg: Security flaws in software can be tough to find. Purposefully planted ones—hidden backdoors created by spies or saboteurs—are often even stealthier. Now imagine a backdoor planted not in an application, or deep in an operating system, but even deeper, in the hardware of the processor that runs a computer. And now imagine that silicon backdoor is invisible not only to the computer’s software, but even to the chip’s designer, who has no idea that it was added by the chip’s manufacturer, likely in some farflung Chinese factory. And that it’s a single component hidden among hundreds of millions or billions. And that each one of those components is less than a thousandth of the width of a human hair.

  In fact, researchers at the University of Michigan haven’t just imagined that computer security nightmare; they’ve built and proved it works. In a study that won the “best paper” award at last week’s IEEE Symposium on Privacy and Security, they detailed the creation of an insidious, microscopic hardware backdoor proof-of-concept. And they showed that by running a series of seemingly innocuous commands on their minutely sabotaged processor, a hacker could reliably trigger a feature of the chip that gives them full access to the operating system. Most disturbingly, they write, that microscopic hardware backdoor wouldn’t be caught by practically any modern method of hardware security analysis, and could be planted by a single employee of a chip factory.

- Named it: "A2" -- because it's an Analog Attack.
- Charge Pump.
- Behavior-based switch -- which would never be found through LOGICAL operation verification.
- Flip the switch and it removes the memory mapper, giving the application unfettered access to the system's entire flat memory model without any hardware access supervision.

Latest Android Security Bulletin / Published yesterday, June 6th
https://source.android.com/security/bulletin/2016-06-01.html

- Made available to Google's Android partners a month earlier, on May 2nd.
  - OTA updates for all recent Nexus versions.
- Google: "The most severe issue is a Critical security vulnerability that could enable remote code execution on an affected device through multiple methods such as email, web browsing, and MMS when processing media files."
- Once again... the Mediaserver module:
  - Critical Mediaserver flaw exposes Android devices to remote code execution; an attacker can send a vulnerable devices a malicious media file that corrupts memory during processing of the file. Because Mediaserver has system- and kernel-level privileges, the resulting code can pretty much do whatever it wants to do.
  - Additionally, an even dozen high-severity issues in Mediaserver permit elevation of privilege that can be exploited by a local malicious app to execute code in the context of the kernel. Google write that an attacker could use this flaw to gain Signature or Signature Or System privileges.

- The WEBM decoder also contained a critical remote code execution flaw. As we know, the open media file format supported by most browsers for video playback. An attacker can exploit this to run code remotely in the context of the Mediaserver process.

- Qualcomm's drivers were also the source of many local privilege escalation vulnerabilities: 5 x Wi-Fi Driver, 2 each in the Video Driver and Sound Driver, 1 x GPU Driver and 1 x Camera. In each case, an attacker could use a malicious application to exploit the flaw and run code in the context of the kernel.

Brad Chacos, senior editor of PC World...

- Fearing forced Windows 10 upgrades, users are disabling critical updates instead
- http://www.pcworld.com/article/3075729/windows/fearing-forced-windows-10-upgrades-users-are-disabling-critical-updates-at-their-own-risk.html
- Yesterday GRC received a panicked note from someone with 117 pending Windows updates.
- Never10 -- 10:40AM, downloads crossed one million.
  - 1,003,622…   (~35k/day)

Mark Zuckerberg’s Twitter and Pinterest accounts hacked, LinkedIn password dump likely to blame

- http://venturebeat.com/2016/06/05/mark-zuckerbergs-twitter-and-pinterests-accounts-hacked-linkedin-password-dump-likely-to-blame/
- A group called "OurMine Team" claimed to have used data from the massive LinkedIn breach to obtain Zuck's password from 2012, which was "dadada" -- and his inactive Twitter account, unused since 2012, was using the same password.

Facebook planning encrypted version of its Messenger bot, sources say

- https://www.theguardian.com/technology/2016/may/31/facebook-messenger-bot-encryption-secure-messaging
- Facebook Messenger is used by ~900 million people.
- The plan is to offer optional end-to-end encryption.
- As with Google's recently announced Allo, the encryption layer will be opt-in since its use would blind Facebook to Messenger content and prevent the use of their planned machine learning features.
- The Guardian noted that last month Google faced some blowback from privacy activists over their opt-in decision... but I think this is entirely appropriate.

5 Things That Give Self-Driving Cars Headaches
http://www.nytimes.com/interactive/2016/06/06/automobiles/autonomous-cars-problems.html
<quote>

- Fully automated cars don't drink and drive, fall asleep at the wheel, text, talk on the phone or put on makeup while driving. With their sensors and processors, they navigate roads without any of these human failings that can result in accidents.

  But there is something self-driving cars do not yet deal with very well – the unexpected. The human brain is still better than any computer at making decisions in the face of sudden, unforeseen events on the road – a child running into the street, a swerving cyclist or a fallen tree limb.

  Here are five situations that, for now at least, often confound self-driving cars and the engineers working on them.

Five tough problems:

- Unpredictable Humans: Bad Drivers
  - Self driving car cannot control the behavior of other drivers.

- Bad Weather: aka "Where did the lines on the road go?
  - Snow, rain, fog, etc. make driving difficult for humans, and it's just a difficult for autonomous systems.
  - Falling snow or rain can also make it difficult for laser sensors to identify obstacles. A large puddle caused by heavy rain may look like blacktop to an autonomous car's sensors.
  - In reports that Google and others have filed with California authorities about their on-road tests of autonomous cars, weather was a prime cause of system failures after which human drivers had to take back control.

- Detours and Rerouted Roads: Variations from pre-existing Digital Maps
  - Google's bubble-shaped self-driving cars rely heavily on highly detailed three-dimensional maps — far more detailed than those in Google Maps — that communicate the location of intersections, stop signs, on-ramps and buildings with the cars' computer systems. Self-driving cars combine these maps with readings from their sensors to find their way around.
  - So... the more you depend upon inherently static pre-existing mapping, the bigger the trouble when something makes those maps wrong.

- Road Discolorations: Is it a pothole or a shadow?
  - Self-driving cars use radar, lasers and high-definition cameras to scan roads for obstacles, and the images they generate are assessed by high-powered processors to identify pedestrians, cyclists and other vehicles. But potholes are tough. They lie below the road surface, not above it. A dark patch in the road ahead could be a pothole. Or an oil spot. Or a puddle. Or even a filled-in pothole.

- ETHICS!: Having to make TOUGH decisions
  - In the midst of busy traffic, a ball bounces into the road, pursued by two running children. If a self-driving car's only options are to hit the children or veer right and strike a telephone pole, potentially injuring or killing the car's occupants, what does it do? Should its computer give priority to the pedestrians or the passengers?

> Engineers are confronting questions like these as they build self-driving technology. When a crash is inevitable and a human is at the wheel, the result is a spontaneous reaction — a decision the driver has to make in a split second. But in a car controlled by algorithms, it is a choice predetermined by a programmer.
>
> This is one of the biggest issues facing the companies working to develop fully autonomous cars, and for now, there’s no concrete solution in sight.

SQRL Update:

SQRL: Identity Lock Protocol

- Jason Figge @jfigge
  @SGgrc What happens if my Sqrl identity is compromised and the 3rd party re-keys my account? How do I recover from this?
- Brian Csipkes @BrianCsipkes
  Steve: love the show. Question about sqrl. This week you were talking about rekeying in sqrl. If I lose control of my private key what prevents the other person from rekeying it before I do? Would they then fully control my accounts?

Miscellany

"The Sequence"

- It's a wonderful perfect little programming environment.

"Human Resource Machine"

- Classic instruction-list programming

The PdDP-8/I

- Now with custom, pre-painted, color-matched PDP-8/I replica rocker switches.
- But... $145? Unbelievable!!
- http://obsolescence.wix.com/obsolescence#!pidp-8/cbie
- http://bit.ly/pdp8kit

HSF:

- 2,677 visitors/day.

SpinRite

Josh in Bartow, Florida
Subject: Long-time listener finally uses SpinRite to recover data
Date: 23 May 2016 18:39:11
:
Hi Steve,

I'm a long time listener to the podcast. I appreciate the work you do each week to keep us safe and informed. Two years ago, I purchased a copy of SpinRite. Like many other listeners, I bought a copy just to support you and everything you do for us without asking for anything in return. I've used SpinRite in maintenance mode a couple of times, but never tried to recover data. Well, until yesterday, that is...

A friend of mine contacted me for help because his computer was acting "weird". He had rebooted the machine and suddenly it was asking him to "Insert Boot Media". The BIOS didn't recognize his hard drive and Windows would not load. He did have a partial backup on an external hard drive; however, it had been a couple of years since the last time he backed everything up. He has two small children and he was worried that he had lost a LOT of precious and irreplaceable photos, along with lots of other files.

I loaned him my copy of SpinRite, recommending that he run it on level 2. This evening, he sent me another message. SpinRite fixed the problem. Windows booted normally and all of his files were intact. My friend is now hard at work backing up all of his files to his external hard drive. I told him to get Carbonite (using the SECURITYNOW offer code, of course). I also told him to buy a copy of SpinRite for himself so he can maintain the drive. He is planning to do both.

Thanks again for making such a great podcast and such a great piece of software.

Baby Monitoring Insecurity

Rapid7
The research presented focuses on the security of retail video baby monitors for a number of reasons. Baby monitors fulfill an intensely personal use case for IoT. They are usually placed near infants and toddlers, are intended to bring peace of mind to new parents, and are marketed as safety devices. By being Internet accessible, they also help connect distant family members with their newest nieces, nephews, and grandchildren, as well as allow parents to check in on their kids when away from home. They are also largely commodity devices, built from general purpose components, using chipsets, firmware, and software found in many other IoT devices.

Video baby monitors make ideal candidates for security exploration; not only are they positioned as safety and security devices (and therefore, should be held to a reasonably high standard for security), but the techniques used in discovering these findings are easily transferable to plenty of other areas of interest. Other products of direct interest to commercial and industrial consumers and security researchers (commercial security systems, home automation systems, on-premise climate control systems) share many of the insecure design and deployment issues found in video baby monitors.

iBaby Labs, Inc. / iBaby M6:

The web site ibabycloud.com has a vulnerability by which any authenticated user to the ibabycloud.com service is able to view the camera information OF ANY OTHER USER, including video recordings, due to a direct object reference vulnerability.

The object ID parameter is eight hexadecimal characters, corresponding to the serial number of the device. This small object ID space enables a trivial enumeration attack, where attackers can quickly brute force the object IDs of all cameras. Once an attacker is able to view an account's details, links provide a filename that is intended to show available "alert" videos that the camera recorded. Using a generic AWS Cloud-Front endpoint found via sniffing iOS app functionality, this URL can have the harvested filename appended and data accessed from the account.

This effectively allows anyone to view videos that were created from that camera stored on the ibabycloud.com service, until those videos are deleted, without any further authentication.

iBaby Labs, Inc. / iBaby M3S

- Runs an open Telnet server with fixed login credentials of "admin"/"admin"
- However, if the camera is behind a NAT router, unless the device's IP and port are mapped, this represent only a local concern. Although UPnP does allow anything inside the network to map the device through to the public Internet.

Philips Electronics N.V.

- The In.Sight B120/37 also runs open local telnet and web services having static well-known passwords.
- But of more concern is the Web Service whose pages contain multiple Cross-Site Scripting (XSS) vulnerabilities which allows any valid account to obtain access to the video streaming of any other account.
- And of still more concern is that all camera streams are proxied though a public cloud provider named Yoics with a public hostname and port. The ports appear to range from 32,000 to 39,000 with the bound port tied to a hostname having the pattern "proxy[1,3-14].yoics.net. Given this small and readily enumerated parameter space, attackers can test for an HTTP 200 response and obtain access to the camera through the public proxy. Once found, administrative privilege is available without authentication of any kind to the web scripts available on the device. And by accessing a streaming URL, a live video and audio stream will be accessible from the camera which appears to remain open on that host/port combination.

Summer Infant Baby Zoom WiFi Monitor & Internet Viewing System

- An authentication bypass allows for the addition of an arbitrary account to any camera, without authentication.
- The web service MySnapCam is used to support the camera's functionality, including account management for access.
- A URL retrievable via an HTTP GET request can be used to add a new user to the camera.
    - https://swifiserv.mysnapcam.com/register/?fn={first_name}&ln={last_name}&email={email}&user-Type=3&userGroup={id} [no authentication required]
- This URL does not require any of the camera's administrators to have a valid session to execute this request, allowing anyone requesting the URL with their details against any camera ID to have access added to that device.
- After a new user is successfully added, an e-mail will then be sent to an e-mail address -- provided by the attacker -- with authentication details for the MySnapCam website and mobile application. Camera administrators are not notified of the new account.
- (Rapid7 also found a web admin privilege escalation... but who cares?)

Lens Laboratories / Lens Peek-a-View
Multiple local open web servers with default "user/user" or "guest/guest" credentials.
(Rapid7 was unable to locate a domain for this vendor.)

Several other devices ship with statically present and open telnet and/or web servers. While this is not a huge concern, per se... if the IoT world continues to be populated by such devices, the likelihood that a scan of the local network will find things, grows... which will strongly incentivize any malware -- meaning any application that you download and run -- to take a quick look around the LAN.