

Security Now! #604 - 03-21-17

Taming Web Ads

This week on Security Now!

Developments in the new windows on old hardware front, Cisco finds a surprise in the Vault7 docs, Ubiquity was caught with the PHPs down, CheckPoint discovered problems in WhatsApp and Telegram, some interesting details about the long running Yahoo breaches, the death of the "ebay Football", the latest amazing IoT insanity, the incredible results of the CanSecWest Pwn2Own competition, a classic "you're doing it wrong" example, Tavis pokes LastPass again, some miscellany and an interesting proposal about controlling web advertising abuse.

Setting Traps for Autonomous Vehicles



Security News

Developments in the "New Windows on Old Hardware" front...

- "Your PC uses a processor that isn't supported on this version of Windows" error when you scan or download Windows updates
- <https://support.microsoft.com/en-us/help/4012982/discusses-an-issue-in-which-you-receive-a-your-pc-uses-a-processor-that>
- Symptoms
 - When you try to scan or download updates through Windows Update, you receive the following error message:
 - Unsupported Hardware
Your PC uses a processor that isn't supported on this version of Windows and you won't receive updates.
 - Additionally, you may see an error message on the Windows Update window that resembles the following:
 - Windows could not search for new updates
An error occurred while checking for new updates for your computer.
Error(s) found:
Code 80240037 Windows Update encountered an unknown error.
- Cause
 - This error occurs because new processor generations require the latest Windows version for support. For example, Windows 10 is the only Windows version that is supported on the following processor generations:
 - Intel seventh (7th)-generation processors
 - AMD "Bristol Ridge"
 - Qualcomm "8996"
 - Because of how this support policy is implemented, Windows 8.1 and Windows 7 devices that have a seventh generation or a later generation processor may no longer be able to scan or download updates through Windows Update or Microsoft Update.
- Resolution
 - We recommend that you upgrade Windows 8.1-based and Windows 7-based computers to Windows 10 if those computers have a processor that is from any of the following generations:
 - Intel seventh (7th)-generation "Intel Core" processor or a later generation
 - AMD seventh (7th)-generation ("Bristol Ridge") processor or a later generation
 - Qualcomm "8996" processor or a later generation
- Press...

- Forbes: "Microsoft Admits Forcing More Users Onto Windows 10"
<https://www.forbes.com/sites/ianmorris/2017/03/17/microsoft-admits-forcing-more-users-onto-windows-10/#51363e8a780f>
- HotHardware: "Microsoft Disables Windows Updates For Ryzen And Kaby Lake PCs Running Windows 7/8.1"
<http://hothardware.com/news/microsoft-disables-windows-update-on-ryzen-kaby-lake-pcs-running-windows-7-windows-8>
- HotHardware: "Microsoft Apparently Ramping Up Heavy-Handed Tactics To Force Windows 10 Migrations"
<http://hothardware.com/news/microsoft-apparently-ramping-up-heavy-handed-tactics-to-force-windows-10-migrations>
- Neowin: "Some new PCs running Windows 7 and 8.1 won't receive further updates"
<https://www.neowin.net/news/some-new-pcs-running-windows-7-and-81-wont-receive-further-updates>
- The Hacker News: "Microsoft Started Blocking Windows 7/8.1 Updates For PCs Running New Processors"
<https://thehackernews.com/2017/03/windows-update-processor.html>
- Processor Microarchitectures & Generations:
 - 2011 - Sandy Bridge
 - 2013 - Haswell
 - (I purchased final Win7 laptop and desktop systems with Haswell.)
 - 2015 - Skylake (6th generation)
 - Microsoft says it will be maintaining a list of approved Skylake systems that are guaranteed to have Windows 7 and 8.1 support through July 17th, 2017 (but not to 2020).
 - 2016 - KabyLake
 - 2017 - Cannonlake

Cisco found a surprise among the leaked CIA Vault7 documents...

- "Cisco IOS and IOS XE Software Cluster Management Protocol Remote Code Execution Vulnerability"
- <https://tools.cisco.com/security/center/content/CiscoSecurityAdvisory/cisco-sa-20170317-cmp>
- SUMMARY
A vulnerability in the Cisco Cluster Management Protocol (CMP) processing code in Cisco IOS and Cisco IOS XE Software could allow an unauthenticated, remote attacker to cause a reload of an affected device or remotely execute code with elevated privileges.

The Cluster Management Protocol utilizes Telnet internally as a signaling and command protocol between cluster members. The vulnerability is due to the combination of two factors:

- The failure to restrict the use of CMP-specific Telnet options only to internal, local communications between cluster members and instead accept and process such options over any Telnet connection to an affected device, and
- The incorrect processing of malformed CMP-specific Telnet options.
- An attacker could exploit this vulnerability by sending malformed CMP-specific Telnet options while establishing a Telnet session with an affected Cisco device configured to accept Telnet connections. An exploit could allow an attacker to execute arbitrary code and obtain full control of the device or cause a reload of the affected device.

Cisco will release software updates that address this vulnerability. There are no workarounds that address this vulnerability.

- Mitigations:
 - This is the high-end IOS-based routers and switches, NOT the consumer Cisco/Linksys boxes.
 - Shutdown TELNET access to the switch
 - -OR- Apply an access control list to allow only other trusted devices to have TELNET access.

Ubiquiti was caught with their PHP down around their ankles.

- Toward the end of November, about four months ago, security researcher Thomas Weber with SEC Consult in Vienna, discovered a command injection flaw in the "pingtest_action.cgi" script. The vulnerability is caused, in part, by some Ubiquiti products employing a 20-year old version of PHP, v2.0.1 from 1997.

The vulnerability can be exploited by luring an attacked user to click on a crafted link or just surf on a malicious website. The whole attack can be performed via a single browser GET-request because there is no Cross Site Request Forgery (CSRF) protection. This allows an attacker to open a port or a reverse shell to connect to the device and is also able to change the device's password since the web service which executed the CGI command runs with root privileges. And even low privileged read-only user accounts, which can be created in the web interface, are able to perform this attack.

- UniFi, EdgeMAX and AmpliFi are not affected.
- Ubiquiti's WiFi "AirOS"-based products are affected.
- Our favored EdgeRouter X (models ER-X and ER-X-SFP) using EdgeOS are not affected.
- There was a lot of back and forth with Ubiquiti. Thomas was patient and kept reaching out. At first Ubiquiti said they were already fixing it, then they said the Proof of Concept samples were not working. Then they said they were, and that it was being fixed. But finally, after several no replies, five days ago Thomas went public with the news -- but he still withheld the PoCs out of respect for their efforts.
- Two days later, on March 18th, patches were available for all affected devices.

- https://www.sec-consult.com/fxdata/secons/prod/temedia/advisories_txt/20170316-0_Ubiquiti_Networks_authenticated_command_injection_v10.txt
- AirOS Vulnerability Issue Update, 3/18/17 - Ubiquiti Networks Community
<https://community.ubnt.com/t5/AirMAX-General-Discussion/AirOS-Vulnerability-Issue-Update-3-18-17/td-p/1869309>

Check Point Discloses Vulnerability that Allowed Hackers to Take over Hundreds of Millions of WhatsApp & Telegram Accounts

- Both Telegram and WhatsApp offer web browser interfaces in addition to native client applications.
- On March 7th, CheckPoint discovered and privately/responsibly reported to both companies a serious web-side attack which was made possible by the way they were both using the browser.
- In both cases, CheckPoint found a means for bypassing each system's file upload authentication, allowing a malicious party to run malicious content in the target victim's browser. This, in turn, provided access to the application's "local store" of sensitive account information.
- A malicious party could send a file (for example, an image) which, when viewed would immediately compromise the user's communications account giving a remote hacker full access to the user's account and all account data, including their profile, personal photos, chat history, contact list, etc. The attacker could then resend the malicious image to all of the compromised user's contacts.
- Both companies responded quickly and repair the vulnerabilities.
- Full details are available in CheckPoint's blog post:
<http://blog.checkpoint.com/2017/03/15/check-point-discloses-vulnerability-whatsapp-telegram/>


Here's how the FBI says Russian hackers stole Yahoo account secrets

- <http://www.cbc.ca/beta/news/technology/russian-yahoo-hackers-indictment-500-million-emails-how-1.4029532>
- The FBI has indicted four men -- three Russians and one Canadian -- with charges of hacking Yahoo, starting back in early 2014, to obtain access to half a billion Yahoo accounts.
- According to the FBI's forensic investigations and subsequent allegations, Yahoo was a persistent target of these attacks.

- A portion the attacks were Spear Phishing, inducing targeted recipients to open malicious documents or click on links taking them to spoofed Phishing websites where, for example, Google Gmail credentials would be requested.
- The attackers also targeted Yahoo's own employees, eventually obtaining their account credentials and leveraging that access to gain access to a Yahoo server in early 2014 and deeper access into Yahoo's corporate network by later that year, in September. They installed their own software to cover the traces of their intrusion and scrub server logs.
- In October 2014 they obtain information about Yahoos "AMT" — Account Management Tool — which Yahoo administrators used to manage and modify information about accounts — user names, recovery email addresses and phone numbers, security questions and answers, and more.
- That information was stored in Yahoo's User Database, or UDB, which, a month later in November, they managed to obtain a backup copy of by early November 2014. This contained information for more than 500 million accounts — providing them access to the account of any user whose password had not been changed since the backup was made.
- Throughout the following two years, 2015 and 2016, the attackers used their access to Yahoo's AMT and the information contained within the stolen UDB to target user accounts of interest.
- They eventually figured out how to mint their own session cookies to bypass the need for username and password account credentials. They initially used Yahoo's own servers for this minting, but by August 2015 they had obtained Yahoo's cookie minting code which allowed them to mint account session cookies on their own machines.
- The indictment alleges that the attackers used these cookies to access the contents of more than 6,500 Yahoo user accounts and later to steal contact information from 30 million Yahoo accounts as part of a spam marketing campaign.
- Spear Phishing of non-Yahoo accounts allowed them to obtain access to at least 80 other accounts, at least 50 of them Gmail.
- The FBI's indictment also notes that the intrusion was powerful enough to allow one of the attackers to modify the results returned by Yahoo searches for a specific erectile dysfunction drug so as to refer to a specific online pharmacy... for which the attacker nobtained a referral fee.
- And throughout all of this... Russian intelligence officials proactively assisted the attackers in hiding their Internet traffic and remaining undetected.

The death of the ebay Football

- @SGgrc ebay is "upgrading" 'a higher level of security' frm token (football) to txt 2 factor :(clue on how secure this is "more convenient"
- Jonathan Bishop (@darthbinkly)
Just got email from ?@eBay?. They're ditching the hardware OTP token and going to texted pins. Ugh. @SGgrc #youredoingitwrong



- (Let's talk about the use of "wrote-only" TOTP (Time based One-Time Password) 6-digit token authenticators.)

This Week in IoT Insanity (You really cannot make this stuff up):

- Standard Innovations, a Canadian company, manufacturers and sells the "We-Vibe"... a bluetooth-enabled, remotely controlled, vibrator.
- It, itself, has no "user-interface" (though some might argue that the entire device is, itself, a user-interface)

- In order to be controlled, it is paired with its companion "We-Connect" App, which allows its user to vary rhythms, patterns and settings — or to give a partner, in the same location or anywhere in the world, control of the device.
- Customers have alleged and complained that the manufacturer was secretly tracking their use, and after a class action lawsuit was filed, the manufacturer has recently reached a \$3.75 million dollar class action settlement with its users following their allegations that the company was collecting data on when and how the sex toy was used.
- NSFW Promotional Video: <https://www.youtube.com/watch?v=4vwYQUo2LSw>
- Slogan: "Play Together... even when you're apart."
- The lawsuit was filed in federal court in Illinois last September, alleging that without customers' knowledge the app was designed to collect information about how often, and with what settings, the vibrator was used.
- Lawyers for the anonymous plaintiff group contended that the app collected users' email addresses to allow the company "to link the usage information to specific customer accounts."
- The suit alleges that customers' email addresses and usage data were transmitted to the company's Canadian servers. When a We-Vibe was remotely linked to a partner, the connection was described as secure, but information was routed through We-Connect and collected.
- The unhappy users allege in their lawsuit that they never agreed to the collection of this data. Standard Innovations maintains that users "consented to the conduct alleged" — but instead of taking the case to court, the company agreed to settle.
- According to court documents an estimated 300,000 people bought Bluetooth-enabled WeVibes.
- Under the terms of the settlement, anyone who bought an app-enabled vibrator can receive up to \$199 dollars; anyone who actually connected it to the app can collect up to \$10,000. The actual amount paid out will depend on how many people file claims; the company estimates people who bought the app will get around \$40, and people who used the app around \$500.
- The vibrators cost between \$119 and \$199, if purchased through the We-Vibe website.
- Standard Innovation also agreed to stop collecting users' email addresses and to update its privacy notice to be clearer about how data is collected.
- In a statement, Standard Innovation called the settlement "fair and reasonable."
- <http://www.businessinsider.com/sex-toy-company-we-vibe-was-fined-c4-million-for-spying-on-its-customers-through-their-smart-vibrators-2017-3>

- <http://www.npr.org/sections/thetwo-way/2017/03/14/520123490/vibrator-maker-to-pay-millions-over-claims-it-secretly-tracked-use>

You Don't Really Need an Anti-Virus App Anymore

- <http://fieldguide.gizmodo.com/you-dont-really-need-an-anti-virus-app-anymore-1793366446>
- Gizmodo: Ten years ago the first thing you needed to load on a brand new computer were anti-virus and malware applications. The internet was a mine field of malicious content that could infect your entire home network with one errant click. Yet things have changed dramatically. Windows has much more robust security options built in, browsers are smarter, and, hopefully, so are users.

Instead of spending your money on a lot of security software that is often as invasive and irritating as the malware it protects against, think long and hard about going with a minimalist security set up and practicing safe browsing techniques. Below are the scant few software software packages you really need, and where to find them.

- A/V is now often increasing the attack surface by interposing less securely written and tested 3rd-party "shimware".
- These days, browsers and OSes are being continually and proactively updated.
- And many OSes now provide built-in Antivirus and Firewall software.

HTTPS Interception Weakens TLS Security | US-CERT

- <https://www.us-cert.gov/ncas/alerts/TA17-075A>
Alert (TA17-075A) / HTTPS Interception Weakens TLS Security
Original release date: March 16, 2017
- Many organizations use HTTPS interception products for several purposes, including detecting malware that uses HTTPS connections to malicious servers. The CERT Coordination Center (CERT/CC) explored the tradeoffs of using HTTPS interception in a blog post called The Risks of SSL Inspection.

Organizations that have performed a risk assessment and determined that HTTPS inspection is a requirement should ensure their HTTPS inspection products are performing correct transport layer security (TLS) certificate validation. Products that do not properly ensure secure TLS communications and do not convey error messages to the user may further weaken the end-to-end protections that HTTPS aims to provide.

SpinRite

- tommyParalyzed (@tommyparalyzed)
@SGgrc Hey Steve! Have USB 500GB HDD. Do I need USB to SATA converter to plug directly into mobo for SpinRite to recognize it? Thanks! :)
- If the drive is plugged into the motherboard when it's booted, the BIOS should "see" the drive, and then SpinRite will too. So you shouldn't need a separate adapter. You can get into the BIOS to verify that it sees the drive. But the drive WILL need to be connected when the system is booted.

More Security News

Zero Day Initiative

- 10th Anniversary of Pwn2Own with the \$1,000,000 USD of prizes.
- As they do every year, the competition order was decided by random drawing in the contest room on the first day of the competition. This year's event features 11 teams of contestants targeting products across four categories - 30 different hacking attempts in total. Each contestant has three attempts within their allotted time to demonstrate their exploit.
- 360 Security (@mj0011sec) targeting Adobe Reader
SUCCESS: The team used a jpeg2000 heap overflow in Adobe Reader, a Windows kernel info leak, and an RCE through an uninitialized buffer in the Windows kernel to take down Adobe Reader. In the process, they have earned themselves \$50,000 USD and 6 points towards Master of Pwn.
- Two researchers targeting Apple Safari with an escalation to root on macOS
PARTIAL SUCCESS: In a partial win, Samuel Groß (@5aelo) and Niklas Baumstark (@_niklasb) earn some style points by leaving a special message on the touch bar of the Mac. They used a use-after-free (UAF) in Safari combined with three logic bugs and a null pointer dereference to exploit Safari and elevate to root in macOS. They earned \$28,000 USD and 9 Master of Pwn points.
- Tencent Security - Team Ether targeting Microsoft Edge
SUCCESS: Tencent Security – Team Ether successfully exploits Microsoft edge through an arbitrary write in Chakra core. They used a logic bug to escape the sandbox and earn themselves \$80,000 and 10 points for Master of Pwn.
- Chaitin Security Research Lab (@ChaitinTech) targeting Ubuntu Desktop
SUCCESS: The Chaitin Security Research Lab (@ChaitinTech) welcomes Ubuntu Linux to Pwn2Own with a Linux kernel heap out-of-bound access. They earned themselves \$15,000 and 3 Master of Pwn points.

- (Later in the first day) Tencent Security - Team Sniper (Keen Lab and PC Mgr) targeting Adobe Reader
SUCCESS: They used an info leak in Reader followed by a UAF (use after free) to get code execution, then they leveraged another UAF in the kernel to gain SYSTEM-level privileges, winning \$25,000 and 6 Master of Pwn points.
- Chaitin Security Research Lab (@ChaitinTech) targeting Apple Safari with an escalation to root on macOS
SUCCESS: They successfully exploited Apple Safari to gain root access on macOS by using a total of six bugs in their exploit chain including an info disclosure in Safari, four different type confusions bugs in the browser, and an a UAF in WindowServer. This earned the team \$35,000 and 11 points towards Master of Pwn.
- 360 Security (@mj0011sec) targeting Adobe Flash with a SYSTEM-level escalation and a virtual machine escape
SUCCESS: They successfully exploited Adobe Flash and elevates to SYSTEM using 4 bugs. They did not complete the VMware escape bonus portion, but what they demonstrated constitutes a win and nets them \$40,000 and 12 Master of Pwn points.
- Tencent Security – Team Sniper (Keen Lab and PC Mgr) successfully exploits Adobe Flash via a UAF and escalates to SYSTEM with a UAF in the Windows kernel. This earned them \$40,000 and 12 points for Master of Pwn.
- Tencent Security – Lance Team successfully exploits Microsoft Edge by using a UAF in Chakra then elevates to SYSTEM by using a UAF in Windows kernel. They earned themselves \$55,000 and 13 Master of Pwn points.
- The Tencent Security - Team Sniper (Keen Lab and PC Mgr) exploits Microsoft Edge with a SYSTEM-level escalation by using a UAF in Chakra and a UAF in the Windows kernel.
- The 360 Security (@mj0011sec) successfully exploits Microsoft Windows with an out-of-bounds bug in the Windows kernel. Nets them \$15,000 and 4 Master of Pwn points.
- The folks from Tencent Security - Team Sniper (Keen Lab and PC Mgr) elevated privileges in Microsoft Windows through an integer overflow in the kernel. This final act of Day Two earned them \$15,000 and 4 points for Master of Pwn.
- The 360 Security (@mj0011sec) successfully elevates privileges on Apple macOS by using an infoleak and race condition in the kernel. In doing so, they garner \$10,000 and 3 more points for Master of Pwn.
- The 360 Security (@mj0011sec) successfully exploited Apple Safari through an integer overflow and escalated to root using a macOS kernel UAF. This garners them \$35,000 and 11 more Master of Pwn points.
- The Chaitin Security Research Lab (@ChaitinTech) succeeds in elevating in macOS by using an infoleak and out-of-bounds bug in the macOS kernel. In doing so, they netted another \$10,000 and 3 more Master of Pwn points.

- The Chaitin Security Research Lab (@ChaitinTech) team finish their Pwn2Own by exploiting Firefox with an integer overflow and escalating privileges through uninitialized buffer in the Windows kernel.
- Tencent Security - Team Sniper (Keen Lab and PC Mgr) exploits Safari with an integer overflow and escalates to root with an out-of-bounds UAF in WindowServer. This nets them \$35,000 and 11 points for Master of Pwn.
- The 360 Security (@mj011sec) team used a used heap overflow in Microsoft Edge, a type confusion bug in the Windows kernel, and an uninitialized buffer in VMware for a complete virtual machine escape. They more than earn \$105,000 and 27 Master of Pwn points.
- Richard Zhu (fluorescence) leveraged two separate use-after-free (UAF) bugs in Microsoft Edge then escalated to SYSTEM using a buffer overflow in the Windows kernel. The garnered him \$55,000 and 14 points towards Master of Pwn.
- Tencent Security - Team Sniper (Keen Lab and PC Mgr) used a three-bug chain to win the Virtual Machines Escapes (Guest-to-Host) category with a VMWare Workstation exploit. They used a Windows kernel UAF, a VMware info leak and an uninitialized VMware buffer to go guest-to-host. This garnered them \$100,000 and 13 points for Master of Pwn.

Insecure login at: Oil And Gas International

- Firefox gets complaint for labeling unencrypted login page insecure
- Sorry! That's a feature not a bug.
 - <https://arstechnica.com/security/2017/03/firefox-gets-complaint-for-labeling-unencrypted-login-page-insecure/>
- Notes:
 - Source page is HTTP.
 - Post destination does not change that:
 - `<form name="frmLogin" method="post" action="login.aspx" id="frmLogin">`
 - Convenience:
 - "Forgot your password? / Click HERE to receive it via email."

Tavis Ormandy finds more trouble in LastPass:

- 4:20pm · 20 Mar 2017 · Twitter Web Client
Oops, new LastPass bug that affects 4.1.42 (Chrome&FF). RCE if you use the "Binary Component", otherwise can steal pwds. Full report on way
- Tavis Ormandy @taviso
@taviso I have a full exploit working without any prompts on Windows, could be made to work on other platforms. Sent details to LastPass.
- <https://bugs.chromium.org/p/project-zero/issues/detail?id=1209>

- Tavis found a reference to one explicit lastpass subdomain:
["https://1min-ui-prod.service.lastpass.com/*"](https://1min-ui-prod.service.lastpass.com/*)
 ... which could be abused to proxy (send) unauthenticated windows messages to the Lastpass extension. This allows complete access to internal privileged LastPass RPC (Remote Procedure Call) commands.
- There are hundreds of internal LastPass RPCs, but the obviously bad ones are things such as copying and filling in passwords (copypass, fillform, etc). If the LastPass binary component is installed (<https://lastpass.com/support.php?cmd=showfaq&id=5576>), then calls such as "openattach" could be used to run arbitrary code.
- LastPass: The issue reported by Tavis Ormandy has been resolved. We will provide additional details on our blog soon.
- Our takeaway is that security is so difficult that highly complex systems need many smart eyeballs checking the work.

A dialog with the WiFi MAC researchers...

- (The US Naval Academy folks who took the trouble to look carefully at the 802.11 Ethernet WiFi MAC addresses that mobile devices were using when probing for hotspots.)
- Lucas: @SGgrc Thanks for your coverage of our research about MAC address Randomization. You provided the best summary over other media outlets.
- Steve: Thanks very much for your note. I'm glad I got the details right. Your work provided a perfect real world platform for helping to explain the details of what's going on, and why this stuff is so difficult to really get right. And, as I said on the podcast, more importantly, your work will doubtless help to drive the improvements our industry needs.
- Lucas: I want to add that I started to listen to your podcast about a year ago and dreamed of having something that made the show. Furthermore when we published I said to one of my professors "How cool would it be if Steve Gibson picked this up on Security Now." We were both ecstatic when we got caught up on episodes and heard the mention of our research. Thanks again for your coverage!

TweetBacks:

- RealBrainTraining (@RealBrainTrainr)
 @SGgrc If an encrypted file is open when my comp enters sleep the file is open when it wakes. Key in RAM is written to drive, Security risk?

 (Discuss security perimeters.)
- Brian Brady (@CSUBrian)
 @SGgrc will the introduction of 3D XPoint reduce the practicality of memory hard key

derivation on memory constrained devices (mobile)

<https://arstechnica.com/information-technology/2017/03/intels-first-optane-ssd-375gb-th-at-you-can-also-use-as-ram/>

- Justin Wayne (@jcwayne)
@SGgrc re: Let's Spoof - The best, last line of defense against phishing is a password manager. If LastPass doesn't fill I'm extra cautious.
- Q: @SGgrc What version of Conway's Game of Life do you play? Several different versions on App Store.
A: Hey, Brian. I got all of that out of my system back in the 70's. So I haven't played with any of the recent implementations. But I would DEFINITELY encourage those who are unfamiliar with LIFE to spend some time playing. It's truly fun and fascinating.
- King Authur (@White_Excalibur)
@SGgrc am thinking about purchasing a PC from eBay. Could the BIOS have malware after wiping the drive? Can I scan?
- Equalizer (@GuidingFuture)
@SGgrc In the Netherlands we seem to be getting more and more need, not yet in politics so much, but amongst the people, for more and easier ways to electronically vote. Reliably. Securely. Anonymously. (is that a word?). Couldn't something like SQRL be modified 2 that end?

(Discuss SQRL's 3rd-party identity assertion.)

Errata

- Ethernet & Bob Metcalf at 3Com not 3M!

Miscellany

- Square It
 - <http://squareit.io/>
- The Expanse has been renewed for a third season.

Taming the misbehavior of web ads

Steve Meyers / PHP and MySQL scalability / Utah Open Source Foundation

<https://github.com/stevecoug>

O'Reilly's OSCon Speaker: Database optimization for web developers

Steve Meyers has worked as a PHP and MySQL scalability expert for the last 15 years at such companies as Omniture (now part of Adobe), Spark Networks (owner of JDate), and CrimeReports. He now runs some of the largest independent online communities of college sports fans. When he's not too busy with all of that, Steve runs the Ski PHP Conference, assists with the OpenWest Conference, is a core team member of the Utah Open Source Foundation, and runs the Provo Linux User Group.

<https://medium.com/@stevecoug/how-to-fix-the-problem-with-malicious-and-misbehaving-ads-6c175df055ff#.c9yeipamo>

How to fix the problem with malicious and misbehaving ads

In ad circles, I'm known as a "publisher". In simple terms, that means that I serve ads on my site, and the ad networks pay me for serving their ads.

This is a wonderful thing. In 2011, I was able to quit my day job in order to focus on my website. I make enough money now that I don't need a "real job" to pay the bills. I love that advertisers made it possible for me to work for myself.

I also hate advertisers. Unlike many "publisher" sites, my users tend to stay for a while. They don't just show up from a Google search, and then go their merry way. Many of them view 50+ pages per day. I've set up my AdSense settings so that I will inflict the least annoyance possible on my users. This includes not allowing ads that autoplay video with audio, ads that expand past their allotted size, ads that popup, ads that popunder, or anything else likely to annoy my users. If the ads get bad enough, my users won't come back, or (more likely) they'll turn on an ad blocker.

Advertisers and ad networks don't care about any of that. They're just out to make sure their ads get the most visibility. I use Google to serve most of my ads, and they make it difficult for me to block ads that are misbehaving. They have tools that supposedly help with finding and blocking ads, but the tools are rarely useful.

For the last two months, I've been trying to block a rash of ads that are driving my users away. They autoplay video with audio, which is bad enough, but in addition to that they force the focus of the window to the ad. If you try to scroll past the ad, they scroll it back. If you try to click on a form input, they refocus on the ad. Since my site is a message board, that's kind of a problem, as it keeps my users from being able to post messages.

You may think that only shady brands would resort to using ads that are so blatantly misbehaving. The ads I've been trying to block are from Ford, Arby's, and Whataburger. These

are national brands, not no-name companies that don't know any better.

And those are the "respectable" ads. I haven't even discussed malvertising yet. Some ads are trying to load viruses onto your computer. Others will redirect you to the app store, or another website. With all this poorly behaved advertising being sent to our users, is it any wonder many of them resort to using ad blockers?

So, how do we fix this? The answer is actually pretty simple.

Browser-enforced content restrictions

Imagine that, as a publisher, you could specify that the ads on your page would be unable to do certain things. For example, you could specify that audio, popups, expanders, and redirects are not allowed in a certain DIV tag. The browser would ensure that any script loaded from that DIV tag would have those capabilities disabled. Alternatively, the permissions could be denied on the SCRIPT tag. That would have the same effect.

In order for this to work, it would need to specify what behaviors are not allowed, rather than what is allowed. This would allow for additional behaviors to be supported in the future.

Do the browsers have any motivation to do this? Google's main business is advertising, and they're the dominant browser on the market, so perhaps not. However, I think it makes sense for them to support a standard like this, because it's in their interest to keep people from using ad blockers.

Publishers need to be able to control their reputation

If we want advertising revenue to continue to be a reasonable way for a website to make money, something needs to be done. Publishers must be able to keep malicious and misbehaving ads off of their sites. This solution puts the power in the hands of the publishers, where it belongs.

Are the browsers willing to do it?