

Security Now! #673 - 07-24-18

The Data Transfer Project

This week on Security Now!

This week we examine still another new Spectre processor speculation attack, some news on DRAM hammering attacks and mitigations, the consequences of freely available malware source code, the reemergence of concern over DNS rebinding attacks, Venmo's very public transaction log, more Russian shenanigans, the emergence of flash botnets, Apple continuing move of Chinese data to China, another (the 5th) Cisco secret backdoor found, an optional missing Windows patch from last week, a bit of Firefox news and piece of errata... and then we look at "The Data Transfer Project" which, I think, marks a major step of maturity for our industry.

It looks like you're
sending out a
ballistic missile
alert.



Security News

The SpectreRSB vulnerability

<https://arxiv.org/pdf/1807.07940.pdf>

Researchers from the University of California at Riverside have published a paper detailing yet another brand new attack on Intel, AMD and ARM speculative execution which can bypass all existing recent software and firmware mitigations against speculative execution attacks.

RSB = Return Stack Buffer << Discuss the thread stack vs the Return Stack Buffer

<quote>

In this paper, we introduce a new attack vector Spectre like attacks that are not prevented by deployed defenses. Specifically, the attacks exploit the Return Stack Buffer (RSB) to cause speculative execution of the payload gadget that reads and exposes sensitive information. The RSB is a processor structure used to predict return address by pushing the return address from a call instruction on an internal hardware stack (typically of size 16 entries). When the return is encountered, the processor uses the top of the RSB to predict the return address to support speculation with very high accuracy.

We show that the RSB can be easily manipulated by user code: a call instruction, causes a value to be pushed to the RSB, but the stack can subsequently be manipulated by the user so that the return address no longer matches the RSB.

We describe the behavior of RSB in more detail, showing an RSB based attack that accomplishes the equivalent of Spectre variant 1 through manipulation of the RSB instead of mistraining the branch predictor; we use this scenario to explain the principles of the attack.

Using these techniques, we construct a number of attack vectors including attacks within the same process, attacks on SGX enclaves (Intel's Software Guard Extensions -- the highest form of data protection offered by Intel chips), attacks on the kernel, and attacks across different threads and processes.

SpectreRSB bypasses all published defenses against Spectre, making it a highly dangerous vulnerability.

Current systems are fundamentally insecure unless speculation is disabled. However, we believe that it is possible to design future generations of CPUs that retain speculation but also close speculative leakage channels, for example by keeping speculative data in separate CPU structures than committed data.

RowHammer, RAMpage & ZebRAM

Our friends from the Systems and Network Security Group at VU Amsterdam have been making many troubling discoveries about the effects of pounding on Dynamic RAM memory are the group that first brought us RowHammer.

About a month ago the news broke of another DRAM-based hammering attack, dubbed RAMpage, which affects all Android phones released since 2012, when Google introduced the ION memory manager as part of Android 4.0 (Ice Cream Sandwich).

The researchers updated their previous app used to detect vulnerability to Drammer attacks to also identify if a device is vulnerable to RAMpage. The app is not available on the Play Store and must be downloaded and then side-loaded: <https://vvdveen.com/drammer.apk>

They also produced an open source mitigation for these attacks known as GuardION: <https://github.com/vusec/guardion>

And now these prolific researchers have announced a new technique they call "ZebRAM", stating that it is comprehensive software protection against Rowhammer attacks. No details are currently available beyond the title of their forthcoming paper titled: "ZebRAM: Comprehensive and Compatible Software Protection Against Rowhammer Attacks." has just been accepted at OSDI 2018, the 13th USENIX Symposium on Operating Systems Design and Implementation.

OSDI will be held October 8–10, 2018 in Carlsbad, CA, USA... so we'll know more about ZebRAM later this year.

Source code for the Android ExoBot banking Trojan appears on the Internet.

... and this has security researchers concerned...

In an interview, a spokesman and security researcher with the security firm ThreatFabric stated that Exobot is an unusually potent banking trojan, capable of infecting even smartphones running the latest Android versions, which is something that very few Trojans can do.

He said that "All threat actors have been working on timing injects (overlay attacks) to work on Android 7, 8, and even 9. However Exobot really is something new. The Trojan gets the package name of the foreground app without requiring any additional permissions. This is a bit buggy, still, but works in most cases."

"The interesting part is that no Android permissions are required. All other Android banking trojans families are using the Accessibility or Use Stats permissions to achieve the same goal but that means they require user interaction with the victim. Exobot does not.

As a result of the source code first being offered for sale, then on of its purchasers deciding to post it publicly and widely, not only is Exobot's source code now freely accessible, but it is also of highly effective.

Consequently, in the coming months researchers expect to see a rise in Android malware based upon Exobot's technology which is now freely and widely available. Hackers who have better

and more effective exploit code in their hands will not hesitate to roll it into their own attack codebases.

DNS Rebinding Exposes Half a Billion Devices in the Enterprise

<https://www.armis.com/dns-rebinding-exposes-half-a-billion-iot-devices-in-the-enterprise/>

Vulnerable device manufacturers ¹	Representative manufacturers	Estimated number of vulnerable devices, worldwide ²
87% of switches, routers, and access points	Aruba Avaya Cisco Extreme Netgear	14 million
78% of streaming media players/speakers	Apple Google Roku Sonos	5.1 million
77% of IP phones	Avaya Cisco Dell NEC Polycom	124 million
75% of IP cameras	Axis Communications GoPro Sony Vivotek	160 million
66% of printers	Hewlett Packard Epson Konica Lexmark Xerox	165 million
57% of smart TVs	Roku-integrated Samsung Vizio	28.1 million

©2018 Armis, Inc Research on estimated exposure of enterprise devices by DNS Rebinding



<https://medium.com/@brannondorsey/attacking-private-networks-from-the-internet-with-dns-rebinding-ea7098a2d325>

<https://lifelhacker.com/prevent-dns-rebinding-attacks-by-adjusting-your-router-1827022291>

<https://www.grc.com/dns/operation.htm>

And as for the **OUTER** circles and arcs . . .

The outer circle of the resolver status icon shows what, if any, "DNS rebinding attack protection" the corresponding nameserver provides to its querying clients.

DNS rebinding attacks  utilize DNS to fool a browser's scripting security into believing that local resources, such as the user's own computer or router, are located in the same web domain as the script's source. When this occurs, the browser's "**Same Origin Policy**"  protection is bypassed, giving scripts unrestricted access to the local resource. This allows scripts to do bad things such as change LAN router settings or access any resources and computers on the LAN. (That's not good.)

Security conscious DNS nameservers are able to help block these attacks simply by never returning IP addresses that fall within the ranges of IP addresses commonly used with private LAN networks behind a router or the "Localhost IP" of **127.0.0.1** which computers use to refer to themselves.





GRC's DNS Benchmark tests each nameserver to determine whether it blocks (filters) the return of these reserved private IP addresses — in both IPv4 and IPv6 formats. At the time of this feature's release, only the OpenDNS nameservers can be configured to do this, and then only for IPv4, IPv6 versions of these queries are still able to sneak through. Since there is never any reason to return a private IP address from a public DNS request **all nameservers should block the return of**

private IP addresses. Hopefully, more will in the future.

As shown in the nearby diagram, the outer circle is divided into four quadrants with each quadrant associated with an IP address in non-routable private networks:

- An **EMPTY** arc (see the 127.0.0.1 IP in the sample diagram) indicates that **no filtering** is provided by the nameserver for the associated network IP.
- A **BLUE** arc (see the 192 and 10 network IPs in the sample diagram) indicates that filtering **is provided** for **either** the IPv4 or IPv6 style address, **but not both**, by the nameserver for the associated network IP.
- A **GREEN** arc (see the 172 network IP in the sample diagram) indicates that filtering is provided for both the IPv4 or IPv6 style address by the nameserver for the associated network IP.

 **The best possible protection** is therefore represented by a full, unbroken, **green outer ring** signifying that all four network IP ranges are being blocked in both IPv4 and IPv6 formats. While no nameservers are providing this protection at the time of this new feature's release, it is our hope that, with time, many nameservers will be updated to do so. No new programming is required to provide this feature. It is simply a matter of updating the nameserver's configuration file.


 Temporary thin black arcs, as shown in the sample to the left, are presented while the detection of each nameserver's rebinding protection is underway. If rebinding protection is proven not to be present the temporary arc will be removed. If either partial or full (both IPv4 and IPv6) protection is confirmed, the temporary black arc will be permanently replaced by either a thick **green** or **blue** arc for each network range.

NOTE: If you would like to learn more about the consequences and prevention of DNS Rebinding attacks, this was the topic of our [Security Now! podcast #260](#). During that episode, Leo and I explained the problem and discussed all of the details of this at some length. [The whole story is available for download](#) in two .mp3 audio sizes and three styles of textual transcripts.

Episode #260 | 05 Aug 2010 | 82 min.

DNS Rebinding

This week, after catching up on all of the post-BlackHat and DefCon conference news, Steve and Leo plow into the detailed depths of "DNS Rebinding." Together they thoroughly explore this significant and fundamental weakness of the Internet's security.



What's the Problem with today's IoT?: The presumption that everything on the local network is trusted and then the convenience of leveraging that assumption.

Perfect case in point: Sonos

Dan Kaminsky first highlighted the danger of DNS rebinding in 2008.

Short-lived DNS resolutions from a clever DNS server.

GRC's DNS Benchmark has, from its inception in 2010, checked all remote DNS servers for the presence of rebinding protection: <https://www.grc.com/dns/benchmark.htm>

Downloads/day: 1,477 -- Total downloads: 4,198,755

To do this, GRC runs a "rebindtest" DNS server: rebindtest.com

```
nslookup net192.rebindtest.com
192.168.0.1
```

Supports:

- net4
- net10
- net127
- net172
- net192

Also supports IPv6 AAAA queries.

OpenDNS does offer a rebinding filtering option.

Some home routers are adding support for filtering out local non-routable addresses

Some users leave their SoHo routers =INTERNAL= web admin logon set to its defaults under the assumption that the only danger is from people on the outside. But DNS rebinding affords a means of JavaScript running on the inside to access LAN devices ON THE INSIDE.

"Venmo" posts all transactions publicly -- and forever -- by default

<https://venmo.com/api/v5/public?limit=1>

<https://venmo.com/api/v5/public>

(Then click on the various "picture" links.)

Wikipedia: "Venmo is a mobile payment service owned by PayPal. It allows users to transfer money to one another using a mobile phone app. It handled \$12 billion in transactions in the first quarter of 2018."

What Wikipedia doesn't mention is that this very popular Paypal Venmo money transfer service is publicly posting its users' transaction records -- by default -- because the Venmo app's default settings are set to "Public" for all users.

Unless users understand this and deliberately change this setting, all the transactions they make are logged and made available to anyone via the Venmo public API.

Data exposed via this API includes the first and last name of the sender and recipient, Venmo avatars (actually, the photo of the user), the date of the transaction, a comment regarding the transaction, transaction types, and more.

"Hang Do Thi Duc", the privacy advocate who re-discovered Venmo's publicly open API queried and downloaded the data on all of the company's 2017 public transactions —207,984,218, in total. The website he set up: <https://publicbydefault.fyi/> highlights some of what he found.

A significant amount of information can be gleaned from assembling sufficient transaction histories for individuals. He built-up a few historical profiles for some of Venmo's customers. He tracked transactions related to a cannabis reseller, a corn dealer, a family, a few random couples, and the story of a woman with 2,033 Venmo transactions.

Twitter: "Very Stable Genius" @r3cgm

Replying to @SGgrc

"Doesn't make much sense for us old fogies but the whole point of the service is to socialize your transactions, and to a particular set of people that's appealing. Private transactions are antithetical to the service itself. We have to be careful about how we frame privacy here."

It's one thing to see them go by in the app in a feed, to be participating in a public arena for the moment. But I wonder whether those who are doing this appreciate that this is a permanent record in a database that presumably survives forever.

Back in October of 2016 Dan Gorelick brought this to light in his "Scraping Venmo" posting where he examined and reverse engineered their not-very-secret API:

<http://danielgorelick.com/scraping-venmo/>

So perhaps "Very Stable Genius" is right... those of us who think this is worrisome are just old fogies who are not very hip and "with it."

The Russians are (still) coming...

Following on our topic from last week about the DOJ's indictment of those 12 Russian agents I wanted to note that during last week's Aspen Security Forum, Tom Burt, Microsoft's vice president for customer security said that earlier this year they had discovered that a fake Microsoft domain had been established by Russia as the landing page for phishing attacks.

Microsoft said it detected and helped the US government to block Russian hacking attempts against at least three congressional candidates this year. Although Microsoft declined to name the targets, they said the three candidates were "people who, because of their positions, might have been interesting targets from an espionage standpoint as well as an election disruption standpoint."

According to Microsoft, the Russian hackers targeted the candidates' staffers with phishing attacks, redirecting them to a fake Microsoft website, in an attempt to steal their credentials.

Tom Burt said that Microsoft "discovered that these [fake domains] were being registered by an activity group that at Microsoft we call Strontium...that's known as Fancy Bear or APT 28."

The emergence of "Flash Botnets"

https://twitter.com/ankit_anubhav/status/1019647993547550720

Ankit Anubhav @ankit_anubhav

Just in : IoT hacker identifying himself as "Anarchy" has claimed to hack about 18000+ Huawei routers. [In 24 hours] The vulnerability is 2017-17215, leaked last Christmas & used in Satori.

He also takes responsibility for massive uptick in Huawei scanning now as seen in @360Netlab scanmon.

The attacker Anarchy has shared a list of infected victim IPs which at that point, I am not making public for obvious reasons.

The motives are not clear as the attacker only told he is doing this "to make the biggest baddest botnet in town" . Probably DDoS.

Its painfully hilarious how attackers can construct big bot armies with known vulns.

<https://blog.newskysecurity.com/huawei-router-exploit-involved-in-satori-and-brickerbot-given-away-for-free-on-christmas-by-ac52fe5e4516>

CVE-2017-17215 is a well-known exploit that has been abused by at least two versions of the Satori botnet and also by BricketBot as well as many of the smaller Mirai-based offshoots. The code has been published online, so it's there for any would-be bot herders to become in-fact bot herders.

And this hacker, Anarchy is not finished. The botnet author told Anubhav that he plans to target CVE-2014-8361, a vulnerability in Realtek routers exploitable via port 52869.

"Testing has already started for the Realtek exploit during the night," Anubhav told Bleeping Computer in a private conversation today. [Update: Both Rapid7 and Greynoise are confirming that scans for Realtek have gone through the roof today.]

So we now live in a world where it's possible to assemble a large massively destructive DDoS botnet in less than a day.

Apple and China

Last year, as we covered at the time, Apple pulled VPN apps from its Chinese App Store under the direction of the Chinese government.

Then earlier this year, in February, Apple moved the encryption keys and data of its Chinese iCloud users from its US servers to local servers on Chinese soil to comply with the new regulation of the Chinese government.

And now, Apple's Chinese data center partner has transferred iCloud data, belonging to 130 million China-based users, to a cloud storage service managed by a state-owned mobile telecom

provider.

Not surprisingly, all of these actions have raising concerns about the privacy of this data.

However, as we've noted while we've been watching this transpire, Apple users in China may, at least for the time being, prevent their data from being stored in servers located in China simply by selecting a different country when setting up their iCloud account.

It is unclear whether an existing user changing this setting would cause their data to be relocated to Apple's US-based data centers and permanently deleting it from Chinese servers.

Cisco has discovered and removed another undocumented backdoor

Last week Cisco released 25 security updates, among them a critical patch with a severity rating of 9.8 out of 10.

This one was for the Cisco Policy Suite which removed an undocumented password which had been built into the "root" account. The previously secret backdoor, tracked as CVE-2018-0375, is significant due to the location of this system within enterprises. The Cisco Policy Suite exists in three editions for Mobile, WiFi, and Broadband Network Gateways which Cisco sells to ISPs and large corporate clients. It allows network admins to set up bandwidth usage policies and subscription plans for customers/employees by tracking individual users, tier traffic, and enforce access policies.

Cisco said there are no workarounds or mitigating factors and that customers need to install the patch to remove the secret password. They say they discovered the undocumented root password during internal security audits and believe that it may have been left behind during software debugging tests ... as they say most of these incidents end up being.

Being the 5th such secret backdoor to be found, for me it's a head shaker. This is no longer seeming like a mistake... as much as a previous policy. But as I've said before, I'm impressed by the way Cisco is behaving now... while apparently under new management!

Just a note about a patch missing from last week, thanks to Bleeping Computer:

<https://support.microsoft.com/en-us/help/4345459/stop-error-0xd1-after-a-race-condition-occurs-in-windows-7-service-pac>

An optional patch which would mostly be of use to people running Windows Server 2008 R2 or Windows Server 2012 R2 (though also affecting Win7 and Win8.1 workstations), was NOT part of the monthly rollout last week and will apparently need to be manually obtained and installed if desired.

Since I =AM= running Server 2008 R2, I plan to apply it when appropriate:

<http://www.catalog.update.microsoft.com/Search.aspx?q=KB4345459>

<http://www.catalog.update.microsoft.com/Search.aspx?q=KB4345424>

From the descriptions it appears that Windows 10 =did= receive these updates as part of the monthly Windows Update process, but that they are optional for the two older versions.

- Addressed issue in which some devices may experience stop error 0xD1 when you run network monitoring workloads.
- Addresses an issue that may cause the restart of the SQL Server service to fail with the error, "Tcp port is already in use".
- Addresses an issue that occurs when an administrator tries to stop the World Wide Web Publishing Service (W3SVC). The W3SVC remains in a "stopping" state, but cannot fully stop or it cannot be restarted.

Firefox auto-play suppression

Just a quick note that Firefox is on track to add its control over auto-playing video with sounds on web pages, joining the Chrome and Edge browsers with Firefox 63, due near the end of October. The default option is "Always Ask" with a pop-up. But users who prefer not to be asked for every site which attempts to auto-play videos can change the default to always "Allow Autoplay" or always "Block Autoplay" ... and thereafter THAT annoying, intrusive and blessedly short-lived practice will bite the dust.

Errata

Herzi @herzi

In SN671 you imply that the PortaPow USB condom is a dumb USB plug that only forwards the power and ground lines. It does a lot more than that as it contains it's own little logic board.

<http://portablepowersupplies.co.uk/home/portapow-data-blocker-usb-adaptor>

- SmartCharge – Built in chip detects the type of device which is connected and swaps between Apple, Universal and Samsung charging specifications. This prevents the blocker from slowing down charging and can increase charging speed if your charger is sending the wrong signal, eg charging an iPad from a charger which uses the 'universal' charging spec.
- Limitations – This adaptor is not compatible with extra fast charging technologies such as 'Qualcomm Quick Charge' 'Samsung adaptive fast charge' as these require data transfer to be enabled. Your device will still charge at a high speed without these. Some satnavs/dashcams use a proprietary signal and will still try to enter sync mode unless their own charger is used. Some car USB sockets do not provide enough power to charge a device, so a dedicated USB car charger must be used.

SpinRite

Site licensing

The Data Transfer Project

<https://datatransferproject.dev/>

<https://datatransferproject.dev/dtp-overview.pdf>

<https://github.com/google/data-transfer-project>

Facebook, Google, Microsoft, Twitter

The Data Transfer Project was formed in 2017 to create an open-source, service-to-service data portability platform so that all individuals across the web could easily move their data between online service providers whenever they want.

The contributors to the Data Transfer Project believe portability and interoperability are central to innovation. Making it easier for individuals to choose among services facilitates competition, empowers individuals to try new services and enables them to choose the offering that best suits their needs.

What is the Data Transfer Project?

Data Transfer Project (DTP) is a collaboration of organizations committed to building a common framework with open-source code that can connect any two online service providers, enabling a seamless, direct, user initiated portability of data between the two platforms.

Individuals should be able to easily transfer their files and data directly between online service providers.

Data Transfer Project (DTP) extends data portability beyond downloading a copy of your data from your service provider, to providing consumers the ability to directly transfer data in and out of any participating provider.

Data Transfer Project is an open source initiative to encourage participation of as many Providers as possible. DTP will enhance the data portability ecosystem by reducing the infrastructure burden on both service providers and users which should in turn increase the number of services offering portability. The protocols and methodology of DTP enable direct, service-to-service data transfer with streamlined engineering work.

How does it work?

The Data Transfer Project uses services' existing APIs and authorization mechanisms to access data. It then uses service specific adapters to transfer that data into a common format, and then back into the new service's API.

DTP comprises three main components:

Data Models are the canonical formats that establish a common understanding of how to transfer data. Adapters provide a method for converting each Provider's proprietary data and authentication formats into a form that is usable by the system. Task Management Library provides the plumbing to power the system. Data Models represent the data when being transferred between two different companies. Ideally each company would use interoperable APIs (e.g. ActivityPub) to allow data to flow between them. However in many cases that is not the case. In those cases there needs to be a way to transfer the data from one company's representation to another company's representation.

Data Models are clustered together, typically by industry grouping, to form Verticals. A Provider could have data in one or more Verticals. Verticals could be photos, email, contacts, or music. Each Vertical has its own set of Data Models that enable seamless transfer of the relevant file types. For example, the Music vertical could have Data Models for music, playlists and videos.

Ideally, a Vertical will have a small number of well-defined and widely-adopted Data Models. In such a situation, the generally accepted standard will be used as the Data Model for that Vertical across companies. This is not currently the case for most Verticals because Data Models have emerged organically in a largely disconnected ecosystem.

One goal of DTP is to encourage organizations to use common Data Models in their systems, which will happen if organizations take importing and exporting data into consideration when initially designing their systems or providing updates. Using a common Data Model will significantly reduce the need for companies to maintain and update proprietary APIs.

Company Specific Adapters:

There are two main kinds of Adapters: Data Adapters and Authentication Adapters. These Adapters exist outside of a Provider's core infrastructure and can be written either by the Provider itself, or by third parties that would like to enable data transfer to and from a Provider.

Data Adapters are pieces of code that translate a given Provider's APIs into Data Models used by DTP. Data Adapters come in pairs: an exporter that translates from the Provider's API into the Data Model, and an importer that translates from the Data Model into the Provider's API.

Authentication Adapters are pieces of code that allow consumers to authenticate their accounts before transferring data out of or into another Provider. OAuth is likely to be the choice for most Providers, however DTP is agnostic to the type of authentication.

Task Management / The rest is just plumbing:

The Task Management Libraries handle background tasks, such as calls between the two relevant Adapters, secure data storage, retry logic, rate limiting, pagination management, failure handling, and individual notifications. DTP has developed a collection of Task Management Libraries as a reference implementation for how to utilize the Adapters to transfer data between two Providers. If preferred, Providers can choose to write their own implementation of the Task Management Libraries that utilize the Data Models and Adapters of DTP.

What are some use cases?

Individuals have many reasons to transfer data, but we want to highlight a few examples that demonstrate the additional value of service-to-service portability.

- A user discovers a new photo printing service offering beautiful and innovative photo book formats, but their photos are stored in their social media account. With the Data Transfer Project, they could visit a website or app offered by the photo printing service and initiate a transfer directly from their social media platform to the photo book service.
- A user doesn't agree with the privacy policy of their music service. They want to stop using it immediately, but don't want to lose the playlists they have created. Using this open-source software, they could use the export functionality of the original Provider to save a copy of their playlists to the cloud. This enables them to import the lists to a new Provider, or multiple Providers, once they decide on a new service.
- A large company is getting requests from customers who would like to import data from a legacy Provider that is going out of business. The legacy Provider has limited options for letting customers move their data. The large company writes an Adapter for the legacy Provider's Application Program Interfaces (APIs) that permits users to transfer data to their service, also benefiting other Providers that handle the same data type.
- A user in a low bandwidth area has been working with an architect on drawings and graphics for a new house. At the end of the project, they both want to transfer all the files from a shared storage system to the user's cloud storage drive. They go to the cloud storage Data Transfer Project User Interface (UI) and move hundreds of large files directly, without straining their bandwidth.
- An industry association for supermarkets wants to allow customers to transfer their loyalty card data from one member grocer to another, so they can get coupons based on buying habits between stores. The Association would do this by hosting an industry-specific Host Platform of DTP.

The innovation in each of these examples lies behind the scenes: Data Transfer Project makes it easy for Providers to allow their customers to interact with their data in ways their customers would expect. In most cases, the direct-data transfer experience will be branded and managed by the receiving Provider, and the customer wouldn't need to see DTP branding or infrastructure at all.

Why do we need DTP?

Users should be in control of their data on the web, part of this is the ability to move their data. Currently users can download a copy of their data from most services, but that is only half the battle in terms of moving their data. DTP aims make move data between providers significantly easier for users.

<https://datatransferproject.dev/dtp-overview.pdf>

25-page Overview PDF