

CacheOut and SGAxe: How SGX Fails in Practice



Stephan van Schaik, Marina Minkin, Andrew Kwong Daniel Genkin and Yuval Yarom







Recap (2018)



Disclosed in 2018: speculative- and transient-execution attacks



Basic cache hierarchy



Meltdown and Foreshadow leak data from caches





Prompted wide mitigation effort





No longer possible to leak from caches

Recap (2019) - MDS Attacks



Turns out there are more buffers

Recap (2019) - MDS Attacks



Disclosed in 2019: Micro-architectural Data Sampling

Recap (2019) - MDS Attacks



MDS targets internal CPU buffers

- Access other people's data through faulting or assisting loads
- Across processes, VMs and Intel SGX

CSO The second Meltdown: New Intel CPU attacks leak secrets



MDS samples data passing through these buffers



MDS attacks are like drinking from a fire hose



You just get whatever data is in flight!



Hence the name: "Micro-architectural Data Sampling"



No control over what you are getting

Mitigating MDS



Mitigation: verw to flush the internal CPU buffers

Part 2: CacheOut



Mitigating MDS

- Flush internal CPU buffers
- Does not fix the root cause behind MDS!
- You can still leak data from those buffers
- Is flushing buffers sufficient as a mitigation?
- Bonus: can we get some more control?





Observations

"... evicting the previous caches line from the L1d cache ... the processor has to write them back through the memory hierarchy and will do this through the LFB ..." The RIDL paper (https://mdsattacks.com/files/ridl.pdf)

- Evicting the cache forces data into the fill buffer
- ZombieLoad reports 0.1 B/s leakage despite mitigations
- This **residual** leakage is worrisome

New Data Path



Data path for eviction through the LFB

New Data Path



But the LFB is still leaky!

So how do we exploit evictions?









The victim's data is in the fill buffer and cache



However, verw flushes that data from the fill buffer





How do we leak the data from the cache?





The cache is divided into sets





The address determines in which set the data is





The attacker reads addresses mapping to the same set





and fills the cache set with its own data



evicting the victim's data into the fill buffer





While the data gets written back to DRAM



The attacker uses a faulty load to leak the data



The attacker uses a faulty load to leak the data



The attacker uses a faulty load to leak the data

Exploiting CacheOut

- CacheOut leaks at 2.85 KiB/s (rather than 0.1 B/s)
- Targets data at **rest**:
 - Control when we push the data into the LFB
 - Long after the victim accessed the data
- Leak AES and RSA keys
- Leak neural network weights
- Break KASLR and leak kernel data
- Even across VMs, including the hypervisor

Exploiting CacheOut



Dump data from SGX enclaves

SGX based Applications



CacheOut



Software Guard Extensions (SGX) Security Model



Software Guard Extensions (SGX) Security Model



Takeaway: trust is based on the EPID key



Remote Client

AaaS (Attestation as a Service)

<u>@SGAxe_AaaS</u> Will attest to anything tweeted at it

- Signed 100+ quotes within 2 hours
 - Blocked by Github
- After the public release of the paper, key was still trusted for a whole month
- Can't update TCB quickly because SGX users need to install BIOS updates
- Hardcoded MRSIGNER prevents abuse

SGAxe-Bot @SGAxe_AaaS · Jun 9 Replying to @bascule

Your quote "Honest Andrew's Jsed Cars, Certificates, and Genuine Intel SGX Enclaves" has been signed. Your quote and instructions on how to verify it can be found at gist.github.com/1afd7a8efa3e0e.... Visit sgaxe.com for more information.

Your Personal Intel SGX Quote

,	EPID Grou Extended PCE SVN: QE SVN: MRSIGNER MRENCLA CPU SVN: Basename	p ID: Group ID: R: VE: :: ta: This q	0xb5 0x00 0x0a 0x0b SGAx Wher 0x0e 0x00 0xb4 0x92 0xd2 0xfb Hone	0x0b 0x00 0x00 e: Hov 0x0e 0x00 0x0e 0x00 0x47 0xc9 0xc3 0xb6 s And	0x00 0x00 v SGX encla 0x02 0x00 0xcc 0x75 0x98 drew's signed	0x00 0x00 Fails i ves gr 0x05 0x00 0xf8 0x4b 0x9f 0xa8 s sed	n Prac o bad 0x01 0x00 0x21 0xae 0x8f Cars,	tice 0x80 0x00 0x1e 0x28 0x55 0x53 Certif	0x00 0x00 0xcc 0xf9 0x3f 0xf8 icates	0x00 0x00 0xfe 0x8a 0x5a 0x23 5, and Ge	Aaaa @SGAxe @SGAxe	Aaas
e-Bot's gists	× 🎽	Notifications /	Twitter	×	Cacl	neOut			×	GitHut)	×
୯ ଜ		🛈 🔒 https	://github.	.com								
Search or jump	o to		/ P	ull requ	lests	ssues	Marke	tplace	Explo	re		

Your account has been flagged.

SGA)

 \mathbf{C}

Because of that, your profile is hidden from the public. If you believe this is a mistake, contact support to have your account status reviewed.

SGX in Signal

- In May, Signal started prompting users to add pins
- Secure Value Recovery (SVR)
 - Backup contacts in cloud
 - Non-phone # based addressing
- Isn't selling point of signal is that they don't store info about users?
 - Data is encrypted under a key derived from both the user's pin and a random seed
 - Designed so that Signal themselves cannot decrypt the data
 - Problem: relies on SGX









- Encryption key derived from both user's secret pin and a random seed
- User's data still secure even when the user's pin is short and memorable

Leaking C2

- Attacker with access to the Signal servers storing C2:
 - Use CacheOut can recover C2 from L1-cache
 - Requires some side-channel knowledge
- Alternatively: exploit the trust placed in the EPID keys
- Values are replicated across a Raft cluster
 - Raft stores C2 and the random seed as a distributed log
- If cluster's network is compromised (subpoena, coercion, hacked machines, etc):
 - Attacker can then use the stolen EPID keys to forge quote proving that the malicious replica is running on SGX
 - Even though fake replica is not using SGX
 - request Raft network to replicate the entire log
 - Can then be read in plaintext



Summary

- CacheOut breaches SGX's confidentiality, allowing an attacker to masquerade as a legitimate SGX enclave
- Good that Signal is thinking about sidechannels against SGX
 - insert LFENCE before each branch
 - also use retpolines to defend against speculation attacks
 - Doesn't mitigate CacheOut
- SGX cannot be relied upon to limit guessing attempts
 - Signal should still require strong passwords for SVR
- Attack assumes a malicious Signal Server with access to the cluster
 - Cannot do this at home
 - Subpoena, insider threat, hacked machines





Countermeasures

- Dedicated Microcode update released on June 20th
 - Patch your machine
- More details at CacheOutAttack.com

Thank you! Questions?



(ankwong@umich.edu)



(stephvs@umich.edu)