

Microarchitectural Attacks on the Cloud

Cat-and-Mouse Games between Hypervisors and the Microarchitecture

Lukas Giner, Andreas Kogler, Sandro Letter

May 10, 2021


IAIK – Graz University of Technology

What is the purpose of page deduplication?


- Reduce amount of pages

- Reduce amount of pages
- COW-principle

- Reduce amount of pages
- COW-principle
- Efficient way of memory usage



(a) normal condition



(b) with memory de-duplication

Example: Attack to find a specific application in another VM

- Fill pages with random data and data for deduplication

- Fill pages with random data and data for deduplication
- Measuring of time to overwrite pages

- Fill pages with random data and data for deduplication
- Measuring of time to overwrite pages
- Fill pages with application based information

- Fill pages with random data and data for deduplication
- Measuring of time to overwrite pages
- Fill pages with application based information
- Wait for deduplication to take place

- Fill pages with random data and data for deduplication
- Measuring of time to overwrite pages
- Fill pages with application based information
- Wait for deduplication to take place
- Modify pages and measure time to overwrite

So we now know another VM uses this application.
What to do next?

- Attack it with flush + reload

- Attack it with flush + reload
- Because we now have **shared memory**

Are there counter measures?

- Read only pages




- Read only pages
- Obfuscation code




- Read only pages
- Obfuscation code
- ? Memory sanitization ?




In the meanwhile, a Platypus appeared in the cloud!




- Energy measurement interface



- Energy measurement interface
- Slow, but **powerfull** side channel



- Energy measurement interface
- Slow, but **powerfull** side channel
- Xen did not restrict access to the interface



- Energy measurement interface
- Slow, but **powerfull** side channel
- Xen did not restrict access to the interface

Takeaway

Carefully choose what the guests can access.

How does virtualization extend the attack surface?

- Instruction decoding **accesses** iTLB



- Instruction decoding **accesses** iTLB
- Attacker **fills** iTLB with two mappings





- Instruction decoding **accesses** iTLB
- Attacker **fills** iTLB with two mappings
 - Normal 4kB page
 - Huge page with different memory type



- Instruction decoding **accesses** iTLB
- Attacker **fills** iTLB with two mappings
 - Normal 4kB page
 - Huge page with different memory type
- System Lockup, Denial of Service



- Instruction decoding **accesses** iTLB
- Attacker **fills** iTLB with two mappings
 - Normal 4kB page
 - Huge page with different memory type
- System Lockup, Denial of Service
- On some CPUs fixed in software



- Instruction decoding **accesses** iTLB
- Attacker **fills** iTLB with two mappings
 - Normal 4kB page
 - Huge page with different memory type
- System Lockup, Denial of Service
- On some CPUs fixed in software

Takeaway

The VM threat model allows for stronger attacks.

But what about Meltdown?



- Xen PV Hypervisor



- Xen PV Hypervisor
- Shares address space



- Xen PV Hypervisor
- Shares address space
- Classical Meltdown attack **leaks** HV data



- Xen PV Hypervisor
- Shares address space
- Classical Meltdown attack **leaks** HV data
- Proposed solution: migrate to HVM



- Xen PV Hypervisor
- Shares address space
- Classical Meltdown attack **leaks** HV data
- Proposed solution: migrate to HVM

Takeaway

Share as little state as possible. Use HVM if applicable.

So, no shared mappings, are we save now?

- FPU/SIMD context switches are **expensive**



- FPU/SIMD context switches are **expensive**
- Switch when access triggers #NM exception





- FPU/SIMD context switches are **expensive**
- Switch when access triggers $\#NM$ exception
- Suppress fault and **encode** data in side channel



- FPU/SIMD context switches are **expensive**
- Switch when access triggers $\#NM$ exception
- Suppress fault and **encode** data in side channel
- Leak e.g. AES-NI registers



- FPU/SIMD context switches are **expensive**
- Switch when access triggers $\#NM$ exception
- Suppress fault and **encode** data in side channel
- Leak e.g. AES-NI registers
- Disabling Lazy FP switches



- FPU/SIMD context switches are **expensive**
- Switch when access triggers $\#NM$ exception
- Suppress fault and **encode** data in side channel
- Leak e.g. AES-NI registers
- Disabling Lazy FP switches

Takeaway

Lazy mechanics often introduce security risks.

So, don't share mappings and don't be lazy! Are we done now?



Foreshadow-VMM

- Foreshadow by Jo van Bulck et.al + Ofir Weisse et al.

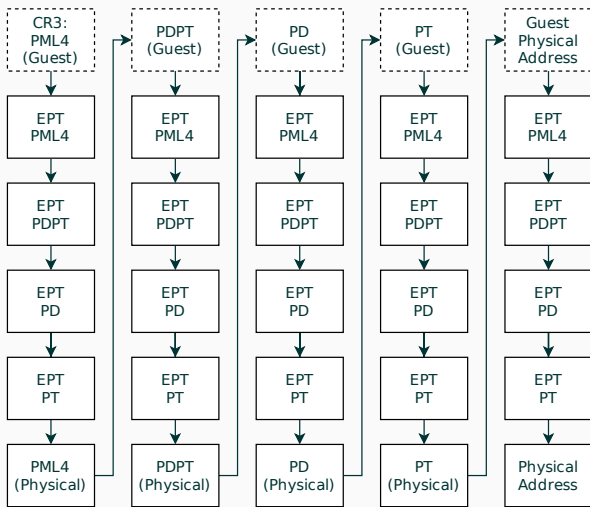


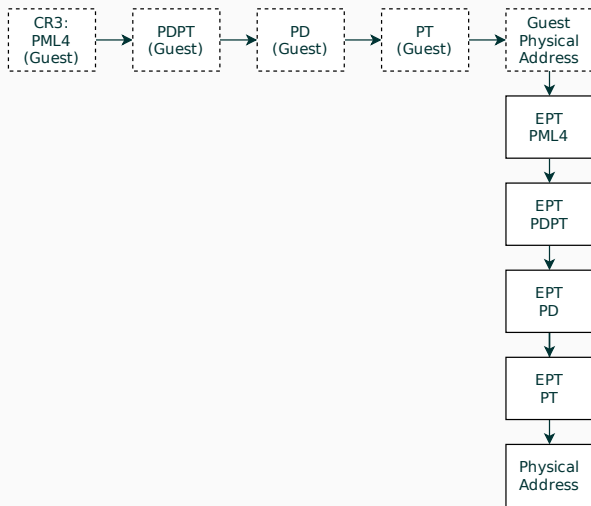
- Foreshadow by Jo van Bulck et.al + Ofir Weisse et al.
- aka L1TF aka Meltdown-P

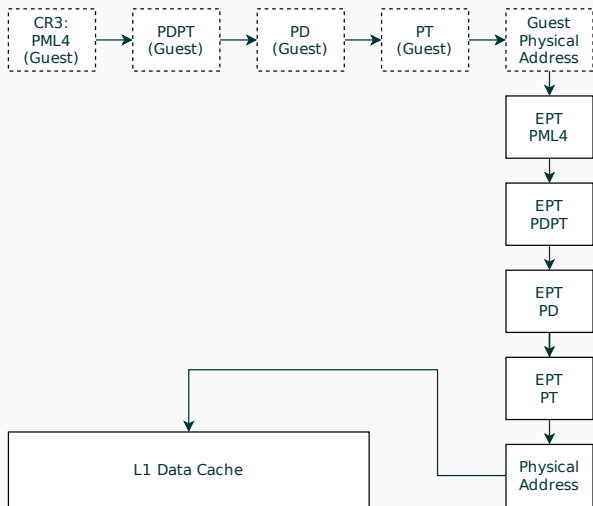


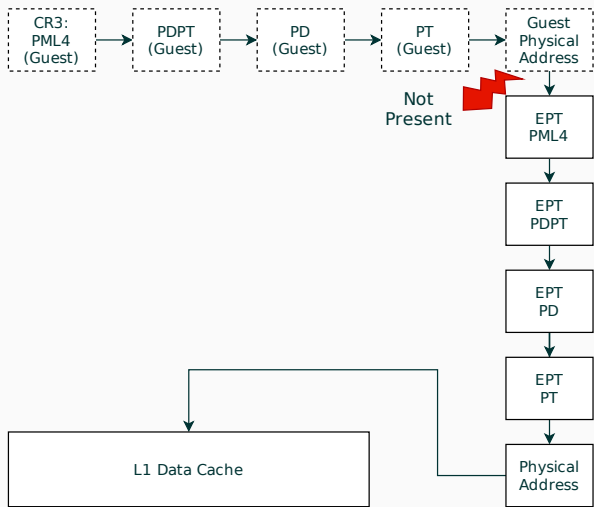
- Foreshadow by Jo van Bulck et.al + Ofir Weisse et al.
- aka L1TF aka Meltdown-P
- (ab)uses the present bit in PTE

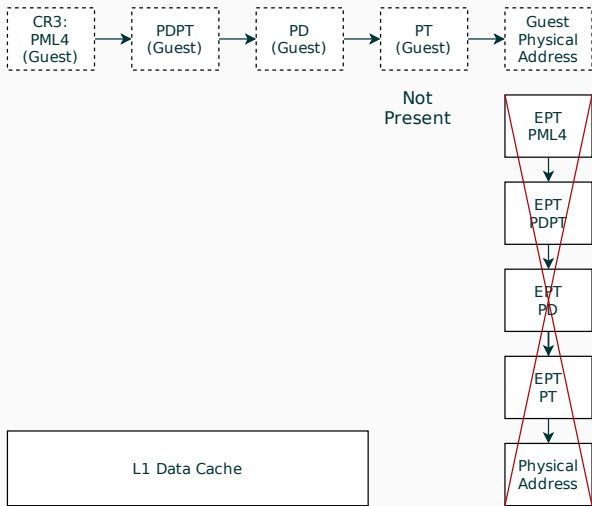


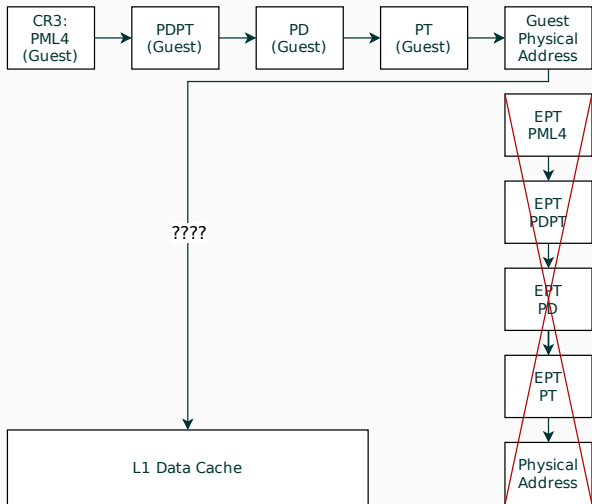














1. VM creates mapping, present = 0



1. VM creates mapping, present = 0
2. CPU stops translation before second level, what now?



1. VM creates mapping, present = 0
2. CPU stops translation before second level, what now?
3. Encode the data from L1 into the cache!



1. VM creates mapping, present = 0
2. CPU stops translation before second level, what now?
3. Encode the data from L1 into the cache!
4. Retrieve with Flush+Reload



1. VM creates mapping, present = 0
2. CPU stops translation before second level, what now?
3. Encode the data from L1 into the cache!
4. Retrieve with Flush+Reload
5. Celebrate



- Fixed in hardware, but until then?



- Fixed in hardware, but until then?
- Flush L1 Cache on entry → no more leakage



- Fixed in hardware, but until then?
- Flush L1 Cache on entry → no more leakage.. right?



- Fixed in hardware, but until then?
- Flush L1 Cache on entry → no more leakage.. right?
- Hyperthreading ruins the day.



- Fixed in hardware, but until then?
- Flush L1 Cache on entry → no more leakage.. right?
- Hyperthreading ruins the day.
- Turn HT off, or only share cores with trusted VMs

- ISA says what should happen..



- ISA says what should happen.. *μ*arch decides how



- ISA says what should happen.. μ arch decides how
- VMM can fix μ arch vulnerabilities..



- ISA says what should happen.. μ arch decides how
- VMM can fix μ arch vulnerabilities.. but also facilitate them!



- ISA says what should happen.. μ arch decides how
- VMM can fix μ arch vulnerabilities.. but also facilitate them!
- VMM has to consider both!

