


Microarchitectural Attacks and Beyond

Daniel Gruss


February 21, 2019


Graz University of Technology

- Bug-free software does not mean safe execution




- Bug-free software does not mean safe execution
- Information leaks due to **underlying hardware**





- Bug-free software does not mean safe execution
- Information leaks due to **underlying hardware**
- **Exploit** leakage through **side-effects**



- Bug-free software does not mean safe execution
- Information leaks due to **underlying hardware**
- **Exploit** leakage through **side-effects**



Power consumption




Execution time




CPU caches







- Instruction Set Architecture (ISA) is an abstract model of a computer (x86, ARMv8, SPARC, ...)



- Instruction Set Architecture (ISA) is an abstract model of a computer (x86, ARMv8, SPARC, ...)
- **Interface** between hardware and software




- Instruction Set Architecture (ISA) is an abstract model of a computer (x86, ARMv8, SPARC, ...)
- **Interface** between hardware and software
- Microarchitecture is an ISA **implementation**




- Instruction Set Architecture (ISA) is an abstract model of a computer (x86, ARMv8, SPARC, ...)
- **Interface** between hardware and software
- Microarchitecture is an ISA **implementation**





- Modern CPUs contain multiple **microarchitectural elements**

- Modern CPUs contain multiple **microarchitectural elements**




Caches and buffer



Predictor





- Modern CPUs contain multiple **microarchitectural elements**




Caches and buffer



Predictor



- **Transparent** for the programmer



- Modern CPUs contain multiple **microarchitectural elements**



Caches and buffer



Predictor



- **Transparent** for the programmer
- Timing optimizations → side-channel leakage



1996



1996



2004



1996



2004



2006



1996



2004



2006



2009



1996



2004



2006



2009



2011



1996



2004



2006



2009



2011



1996



2004



2006



2009



2011



2013



1996



2004



2006



2009



2011



2013





1996



2004



2006



2009



2011



2013





1996



2004



2006




2009



2011



2013



2014



1996



2004



2006




2009



2011



2013



2014





1996



2004



2006




2009



2011



2013



2014





1996



2004



2006




2009




2011



2013



2014





1996



2004



2006




2009




2011



2013



2014



2015



1996



2004



2006




2009




2011




2013




2014




2015




2016




2016



2017



2016





2017





2018


```
printf("%d", i);  
printf("%d", i);
```





```
printf("%d", i);  
printf("%d", i);
```


Cache miss

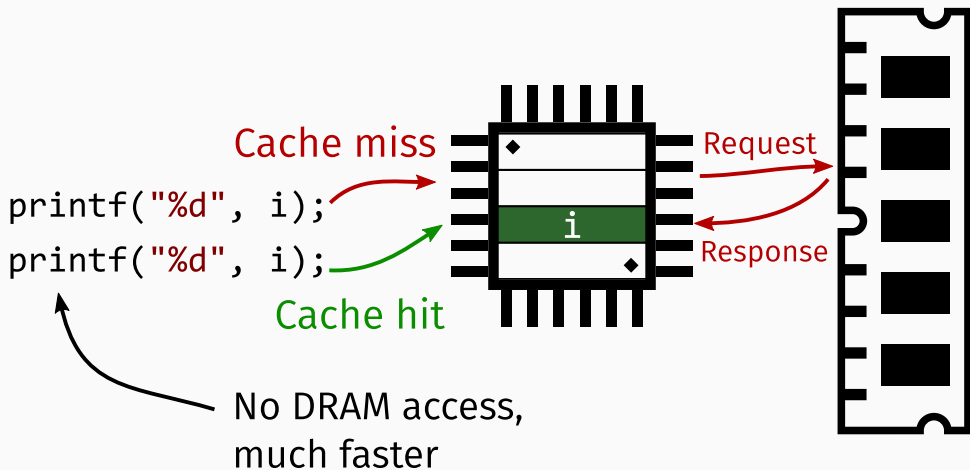













































Terminal

File Edit View Search Terminal Help

```
% sleep 2; ./spy 300 7f05140a4000-7f051417b000 r-xp 0x20000 08:02 26
8050 /usr/lib/x86_64-linux-gnu/gedit/libgedit.so
```

gnome-terminal

gnome-terminal

Terminal

File Edit View Search Terminal Help

```
shark% ./spy
```

gnome-terminal

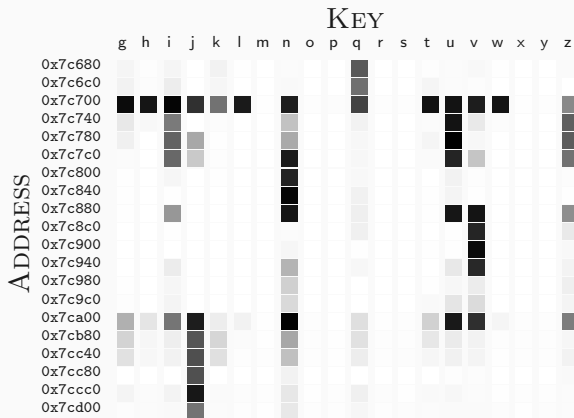
Open +

Untitled Document 1


Save

1

Plain Text Tab Width: 2 Ln 1, Col 1 INS




Raw Prime+Probe trace...¹




¹Michael Schwarz et al. Malware Guard Extension: Using SGX to Conceal Cache Attacks. In: DIMVA. 2017.

...processed with a simple moving average...¹



¹Michael Schwarz et al. Malware Guard Extension: Using SGX to Conceal Cache Attacks. In: DIMVA. 2017.

...allows to clearly see the bits of the exponent¹



¹Michael Schwarz et al. Malware Guard Extension: Using SGX to Conceal Cache Attacks. In: DIMVA. 2017.


YOU CAN'T DO THAT!

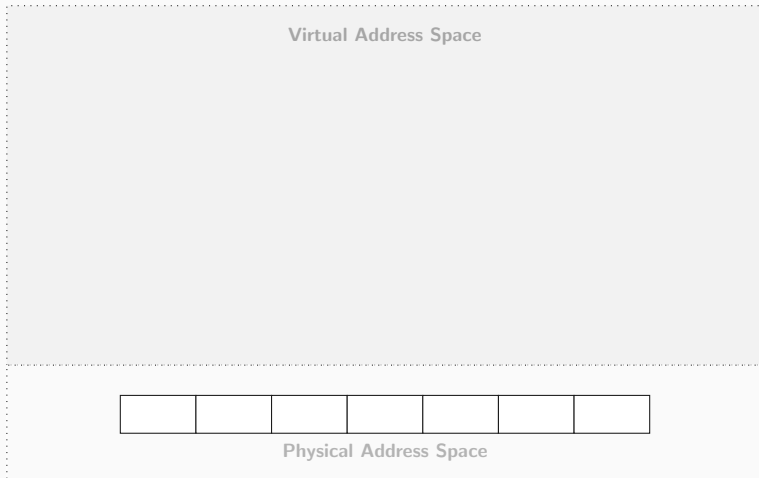


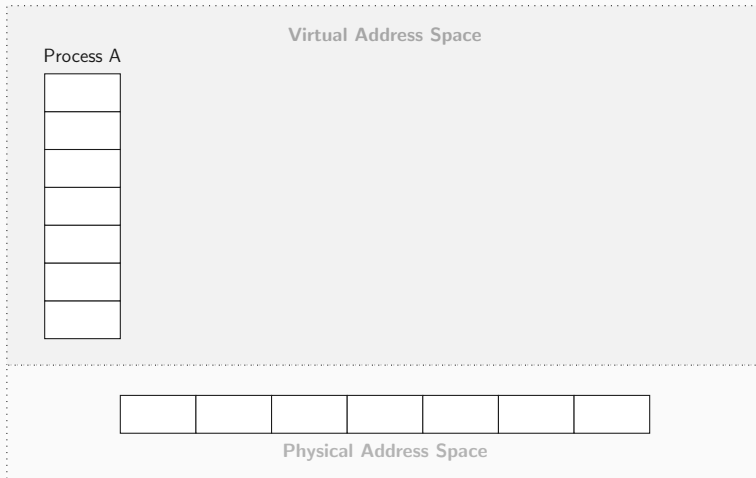
THAT'S AGAINST THE RULES!

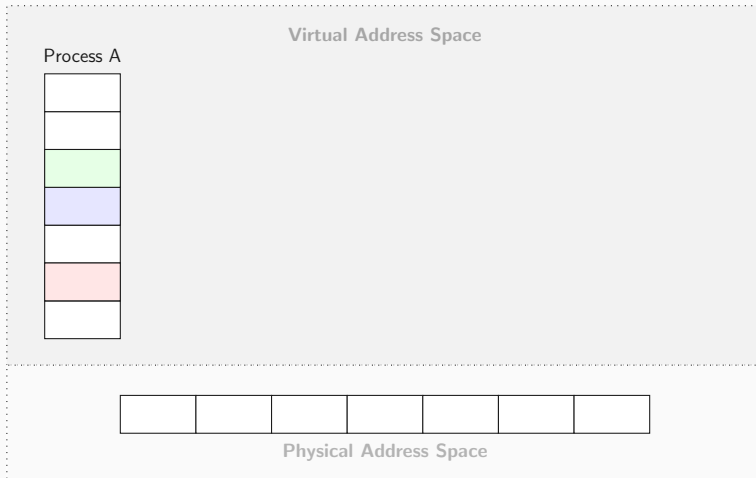
- Files buffered page-wise in “page cache”
- Lower access time for frequently accessed data


- Files buffered page-wise in “page cache”
- Lower access time for frequently accessed data
- Use up all the memory
- Pages are freed on demand
- Deduplicate pages (copy-on-write)



















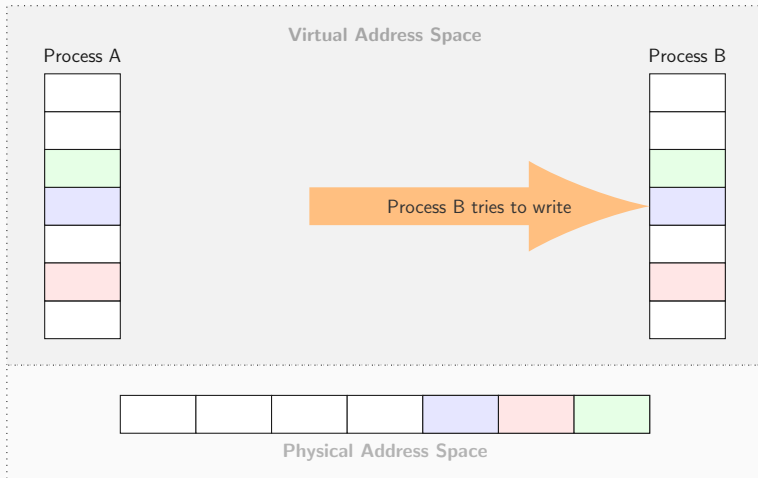












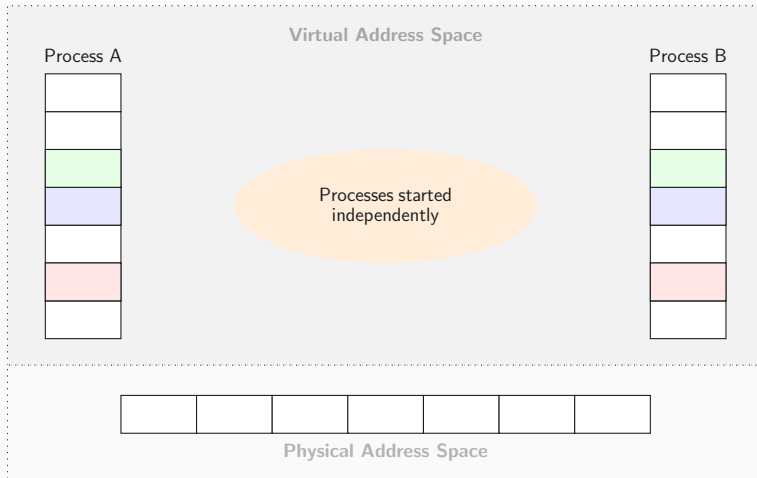









- Regular write access $< 0.2\mu s$
- Write access with copy-on-write pagefault $> 3.0\mu s$
- Clearly distinguishable







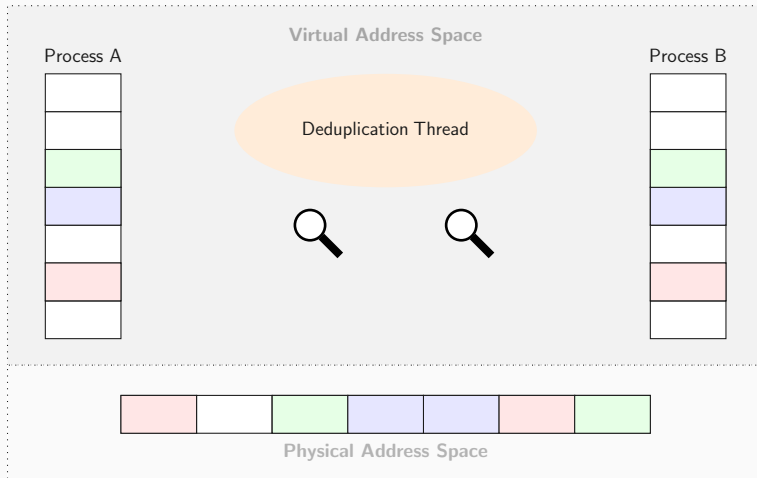




























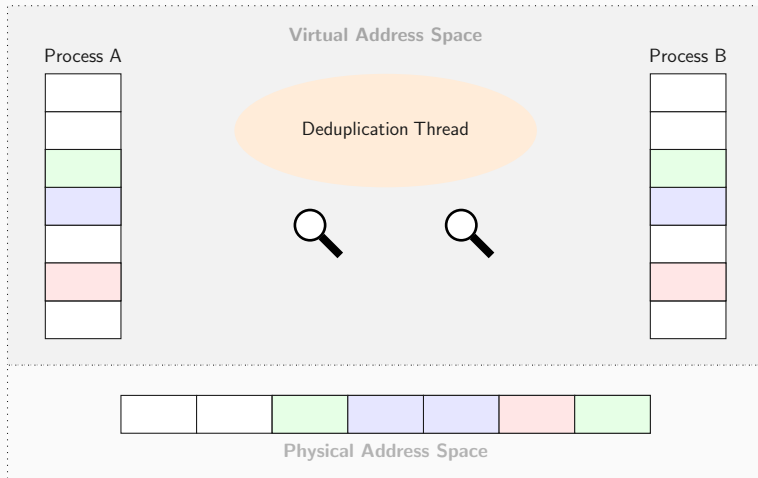









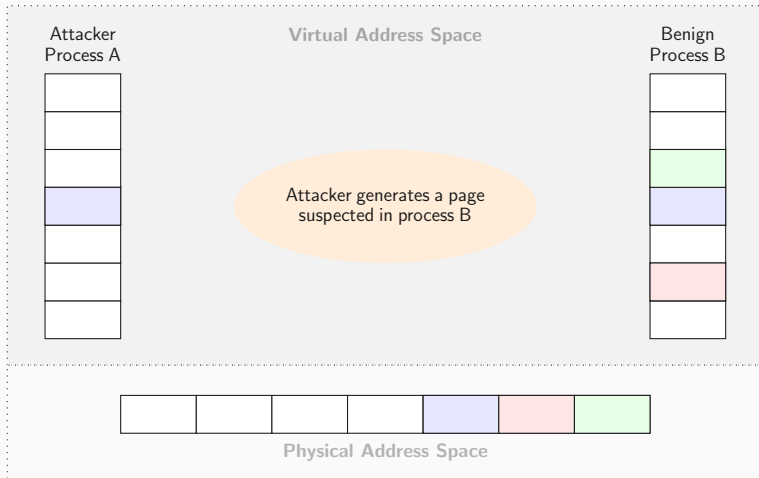






- Deduplication between processes:
 1. in same OS instance (Android, Windows)
 2. in different VMs (KVM, VMWare, ...)


- Deduplication between processes:
 1. in same OS instance (Android, Windows)
 2. in different VMs (KVM, VMWare, ...)
- Code pages, data pages - even kernel pages
- Time until deduplication 2-45 minutes
 - depends on system configuration





















































































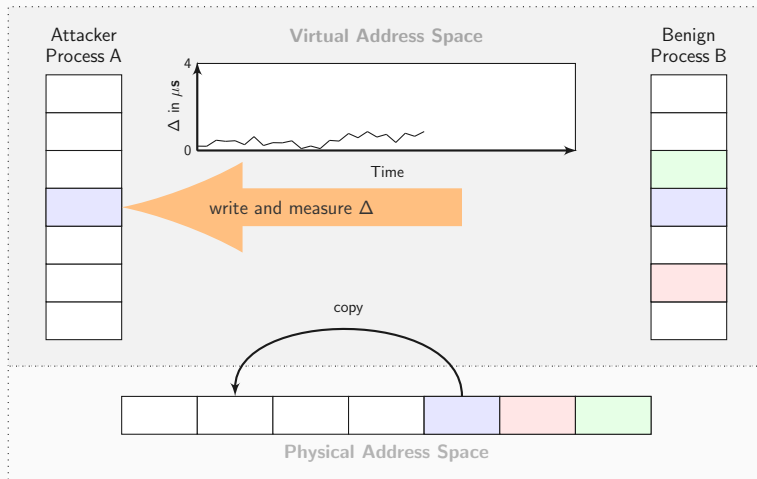








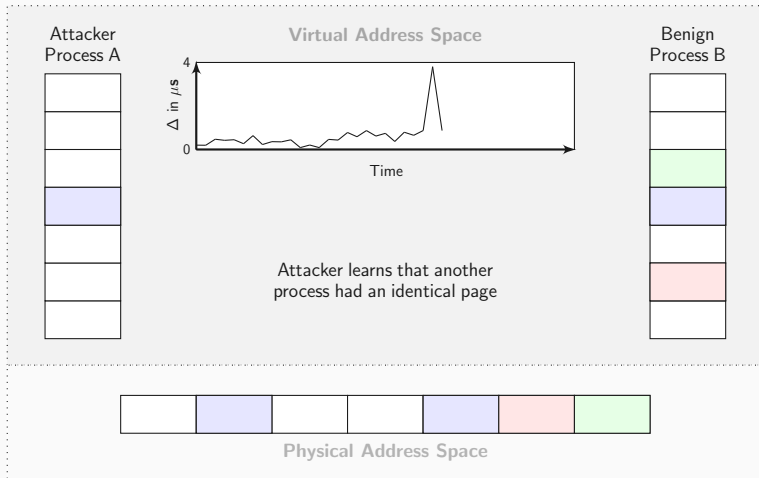


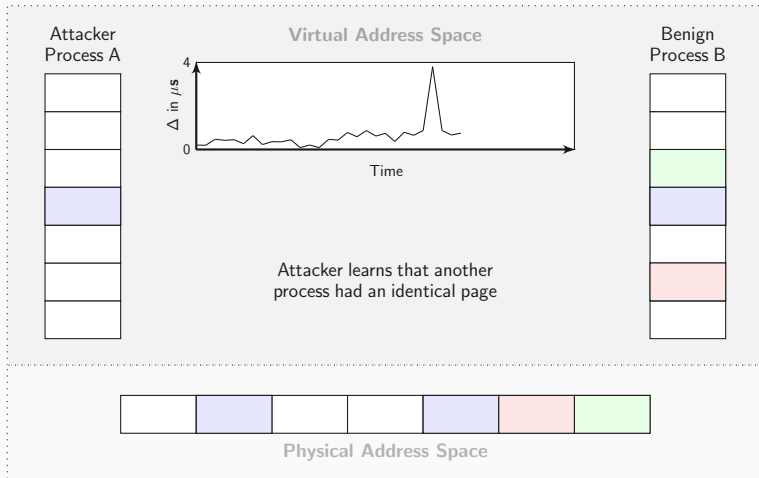






















- Detect binary versions in co-located VMs
 - Detect downloaded image in Firefox under certain conditions
- Attacks on hypervisors
- Native code only


- Detect CSS files and images of opened websites
 - Chrome, Firefox and Internet Explorer
 - Perform the attack in JavaScript
- Attacks on KVM, Windows 8.1 and Android

- Images and CSS files are page-aligned in memory
 - Load them into memory for all websites of interest
 - Detect deduplication
- Malicious ad networks: alternative to tracking pixels?



- No cycle counting (`rdtsc`)
- No access to virtual addresses

- Only require microsecond accuracy
 - `performance.now()` is accurate enough
 - Can even work with millisecond accuracy
 - Accumulate time difference
 - Only possible with enough image/CSS data
- Large typed arrays are allocated page-aligned



- Attacker chosen set of websites
- Load website images and CSS files into arrays
- Reuse HTTP headers of system under attack

JavaScript:

- Reduce timer accuracy?
- Prevent page-aligned arrays?
- Website diversification?
- Prevent control over full pages
 - Every n -th byte not part of JavaScript array

JavaScript:

- Reduce timer accuracy?
- Prevent page-aligned arrays?
- Website diversification?
- Prevent control over full pages
 - Every n -th byte not part of JavaScript array

Generic:

- Disable page deduplication (for writable pages)

- Can we mount an attack without page deduplication?

- Can we mount an attack without page deduplication?
- Shared pages are in the page cache

- Can we mount an attack without page deduplication?
- Shared pages are in the page cache
- Non-shared pages too

```
MINCORE(2)                                Linux Programmer's Manual                                MINCORE(2)
```

NAME [top](#)

`mincore` - determine whether pages are resident in memory

SYNOPSIS [top](#)

```
#include <unistd.h>
#include <sys/mman.h>

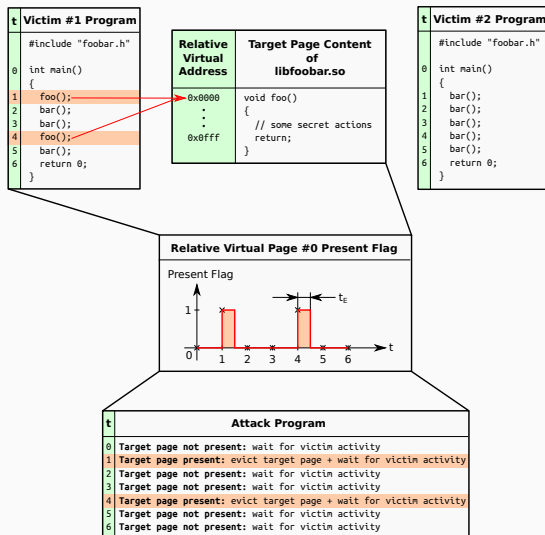
int mincore(void *addr, size_t length, unsigned char *vec);
```

Feature Test Macro Requirements for glibc (see [feature_test_macros\(7\)](#)):

```
mincore():
    Since glibc 2.19:
        _DEFAULT_SOURCE
    Glibc 2.19 and earlier:
        _BSD_SOURCE || _SVID_SOURCE
```

DESCRIPTION [top](#)

`mincore()` returns a vector that indicates whether pages of the calling process's virtual memory are resident in core (RAM), and so will not cause a disk access (page fault) if referenced. The kernel returns residency information about the pages starting at the address `addr`, and continuing for `length` bytes.



- No unprivileged flush system call


- No unprivileged flush system call
- Eviction? Well...

- No unprivileged flush system call
- Eviction? Well...
- Filling your entire memory with garbage takes long, makes your system unstable and laggy


- No unprivileged flush system call
- Eviction? Well...
- Filling your entire memory with garbage takes long, makes your system unstable and laggy
- Got down to around 2–10 seconds with that approach in 2015.

- No unprivileged flush system call
- Eviction? Well...
- Filling your entire memory with garbage takes long, makes your system unstable and laggy
- Got down to around 2–10 seconds with that approach in 2015.
- Idea: use page cache pages for eviction


(1) Start




(2) Evict Page Cache




(3) Access Binary




(4) Evict + Access




(5) Evict + Access



(6) Stop if target reached



- Great advantage over memory exhaustion: only negligible memory footprint



- 2.68s for one eviction, that's still too slow for most use cases

- 2.68s for one eviction, that's still too slow for most use cases
- Idea: Let's build the eviction set more cleverly

- Ideally contains all pages that are currently / most often in the page cache


- Ideally contains all pages that are currently / most often in the page cache
- Eviction will load useful pages into the page cache + not evict these


- Ideally contains all pages that are currently / most often in the page cache
- Eviction will load useful pages into the page cache + not evict these
- Add some pages which are not needed + rarely accessed → evict these first


- Ideally contains all pages that are currently / most often in the page cache
- Eviction will load useful pages into the page cache + not evict these
- Add some pages which are not needed + rarely accessed → evict these first
 - Add memory pressure to reduce the page cache size for faster eviction

- Ideally contains all pages that are currently / most often in the page cache
- Eviction will load useful pages into the page cache + not evict these
- Add some pages which are not needed + rarely accessed → evict these first
 - Add memory pressure to reduce the page cache size for faster eviction

- Ideally contains all pages that are currently / most often in the page cache
- Eviction will load useful pages into the page cache + not evict these
- Add some pages which are not needed + rarely accessed → evict these first
 - Add memory pressure to reduce the page cache size for faster eviction
- 149 ms for one eviction







- `QueryWorkingSetEx` instead of `mincore`

- QueryWorkingSetEx instead of mincore
 - PROCESS_QUERY_LIMITED_INFORMATION? → can even attack non-shared pages (heap, stack, etc.)

- `QueryWorkingSetEx` instead of `mincore`
 - `PROCESS_QUERY_LIMITED_INFORMATION`? → can even attack non-shared pages (heap, stack, etc.)
 - `ShareCount` → requires no permissions + works on shared pages

- QueryWorkingSetEx instead of mincore
 - PROCESS_QUERY_LIMITED_INFORMATION? → can even attack non-shared pages (heap, stack, etc.)
 - ShareCount → requires no permissions + works on shared pages
- Page cache eviction really inefficient

- QueryWorkingSetEx instead of mincore
 - PROCESS_QUERY_LIMITED_INFORMATION? → can even attack non-shared pages (heap, stack, etc.)
 - ShareCount → requires no permissions + works on shared pages
- Page cache eviction really inefficient
- Working set eviction instead:

- `QueryWorkingSetEx` instead of `mincore`
 - `PROCESS_QUERY_LIMITED_INFORMATION`? → can even attack non-shared pages (heap, stack, etc.)
 - `ShareCount` → requires no permissions + works on shared pages
- Page cache eviction really inefficient
- Working set eviction instead:
 - `VirtualUnlock` flushes pages

- `QueryWorkingSetEx` instead of `mincore`
 - `PROCESS_QUERY_LIMITED_INFORMATION`? → can even attack non-shared pages (heap, stack, etc.)
 - `ShareCount` → requires no permissions + works on shared pages
- Page cache eviction really inefficient
- Working set eviction instead:
 - `VirtualUnlock` flushes pages
 - `SetProcessWorkingSetSize` to a minimum (52 KB) + self-eviction in victim

- `QueryWorkingSetEx` instead of `mincore`
 - `PROCESS_QUERY_LIMITED_INFORMATION`? → can even attack non-shared pages (heap, stack, etc.)
 - `ShareCount` → requires no permissions + works on shared pages
- Page cache eviction really inefficient
- Working set eviction instead:
 - `VirtualUnlock` flushes pages
 - `SetProcessWorkingSetSize` to a minimum (52 KB) + self-eviction in victim

- `QueryWorkingSetEx` instead of `mincore`
 - `PROCESS_QUERY_LIMITED_INFORMATION`? → can even attack non-shared pages (heap, stack, etc.)
 - `ShareCount` → requires no permissions + works on shared pages
- Page cache eviction really inefficient
- Working set eviction instead:
 - `VirtualUnlock` flushes pages
 - `SetProcessWorkingSetSize` to a minimum (52 KB) + self-eviction in victim

→ 4.48 ms for one eviction

- Cross-container / cross-sandbox covert channel:

- Cross-container / cross-sandbox covert channel:
 - 7 KB/s on Linux

- Cross-container / cross-sandbox covert channel:
 - 7 KB/s on Linux
 - 273 KB/s on Windows 10

- Cross-container / cross-sandbox covert channel:
 - 7 KB/s on Linux
 - 273 KB/s on Windows 10
- ASLR break: < 1 minute to locate stack, heap, and binary (in other process on Windows)


- Cross-container / cross-sandbox covert channel:
 - 7 KB/s on Linux
 - 273 KB/s on Windows 10
- ASLR break: < 1 minute to locate stack, heap, and binary (in other process on Windows)
- Observe when user plays a video in Firefox, distinguish codecs (and by that sites), when a PNG is rendered, various user input

- Cross-container / cross-sandbox covert channel:
 - 7 KB/s on Linux
 - 273 KB/s on Windows 10
- ASLR break: < 1 minute to locate stack, heap, and binary (in other process on Windows)
- Observe when user plays a video in Firefox, distinguish codecs (and by that sites), when a PNG is rendered, various user input
- UI redress attack


- Cross-container / cross-sandbox covert channel:
 - 7 KB/s on Linux
 - 273 KB/s on Windows 10
- ASLR break: < 1 minute to locate stack, heap, and binary (in other process on Windows)
- Observe when user plays a video in Firefox, distinguish codecs (and by that sites), when a PNG is rendered, various user input
- UI redress attack
- Keystroke timing attack

- Cross-container / cross-sandbox covert channel:
 - 7 KB/s on Linux
 - 273 KB/s on Windows 10
- ASLR break: < 1 minute to locate stack, heap, and binary (in other process on Windows)
- Observe when user plays a video in Firefox, distinguish codecs (and by that sites), when a PNG is rendered, various user input
- UI redress attack
- Keystroke timing attack
- PHP password generation

- Cross-container / cross-sandbox covert channel:
 - 7 KB/s on Linux
 - 273 KB/s on Windows 10
- ASLR break: < 1 minute to locate stack, heap, and binary (in other process on Windows)
- Observe when user plays a video in Firefox, distinguish codecs (and by that sites), when a PNG is rendered, various user input
- UI redress attack
- Keystroke timing attack
- PHP password generation
- Oracle attacks



(Linux)

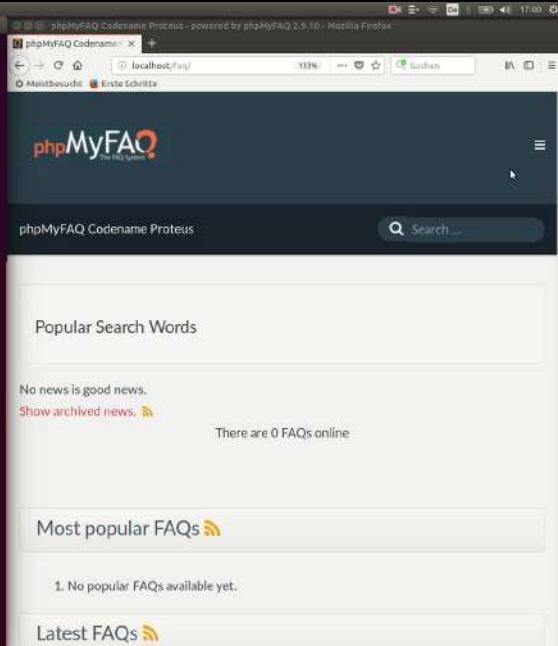


(Windows)



```
erik@erik-Lenovo-Yoga-3-14: ~/data
erik@erik-Lenovo-Yoga-3-14:~/data$ ./eviction -t /usr/lib/policykit-1-
gnome/polkit-gnome-authentication-agent-1 2 ./auth_dialog
[OK] Total usable ram 8246423552
[OK] Target offset 2000...
[OK] File eviction.ram already exists...
[PENDING] Initialising...
[INFO] 1018 libraries are currently resident in memory using approx 30
MB.
```

```
tu@erik-Lenovo-Yoga-3-14: ~/data$ ./eviction -t /usr/sbin/php-  
fpm7.0 441  
[OK] Total usable ram 8246423552  
[OK] Target offset 1b9000...  
[OK] File eviction.ram already exists...  
[PENDING] Initialising...  
[INFO] 1019 libraries are currently resident in memory using  
approx 30 MB.
```





- generic cache attacks also exist without hardware caches



- generic cache attacks also exist without hardware caches
- potentially better suited for malware (hardware-agnostic)



- generic cache attacks also exist without hardware caches
- potentially better suited for malware (hardware-agnostic)
- difficult to mitigate entirely

Microarchitectural Attacks and Beyond

Daniel Gruss

February 21, 2019

Graz University of Technology