

# Side-Channel Lab I

**Michael Schwarz** 

Security Week Graz 2019

Michael Schwarz — Security Week Graz 2019

1













www.tugraz.at





 everyday hardware: servers, workstations, laptops, smartphones...

www.tugraz.at





 everyday hardware: servers, workstations, laptops, smartphones...

• remote side-channel attacks

• safe software infrastructure  $\rightarrow$  no bugs, e.g., Heartbleed

- **safe software** infrastructure  $\rightarrow$  no bugs, e.g., Heartbleed
- does not mean safe execution

Side channels

- safe software infrastructure  $\rightarrow$  no bugs, e.g., Heartbleed
- does not mean safe execution
- information leaks because of the hardware it runs on

- safe software infrastructure  $\rightarrow$  no bugs, e.g., Heartbleed
- does not mean safe execution
- information leaks because of the hardware it runs on
- no "bug" in the sense of a mistake  $\rightarrow$  lots of performance optimizations

- safe software infrastructure  $\rightarrow$  no bugs, e.g., Heartbleed
- does not mean safe execution
- information leaks because of the hardware it runs on
- no "bug" in the sense of a mistake  $\rightarrow$  lots of performance optimizations
- $\rightarrow\,$  crypto and sensitive info., e.g., keystrokes and mouse movements





#### Why targeting the cache?

• shared across cores

#### Why targeting the cache?

- shared across cores
- fast

#### Why targeting the cache?

- shared across cores
- fast
- $\rightarrow$  fast cross-core attacks!



• caches improve performance

- caches improve performance
- SRAM is expensive  $\rightarrow$  small caches

- caches improve performance
- SRAM is expensive  $\rightarrow$  small caches
- different timings for memory accesses

- caches improve performance
- SRAM is expensive  $\rightarrow$  small caches
- different timings for memory accesses
  - data is **cached**  $\rightarrow$  cache hit  $\rightarrow$  **fast**

- caches improve performance
- SRAM is expensive  $\rightarrow$  small caches
- different timings for memory accesses
  - data is **cached**  $\rightarrow$  cache hit  $\rightarrow$  **fast**
  - data is **not cached**  $\rightarrow$  cache miss  $\rightarrow$  **slow**







www.tugraz.at



• L1 and L2 are private

Michael Schwarz — Security Week Graz 2019

www.tugraz.at



- L1 and L2 are private
- last-level cache:

Michael Schwarz — Security Week Graz 2019



- L1 and L2 are private
- last-level cache:
  - divided in slices



- L1 and L2 are private
- last-level cache:
  - divided in slices
  - shared across cores



- L1 and L2 are private
- last-level cache:
  - divided in slices
  - shared across cores
  - inclusive

















www.tugraz.at

## **Inclusive property**



- inclusive LLC: superset of L1 and L2
- data evicted from the LLC is also evicted from L1 and L2

www.tugraz.at

## **Inclusive property**



- inclusive LLC: superset of L1 and L2
- data evicted from the LLC is also evicted from L1 and L2
- a core can evict lines in the private L1 of another core



On current Intel CPUs:

• Registers: 0-1 cycle



On current Intel CPUs:

- Registers: 0-1 cycle
- L1 cache: 4 cycles
www.tugraz.at

On current Intel CPUs:

- Registers: 0-1 cycle
- L1 cache: 4 cycles
- L2 cache: 12 cycles

www.tugraz.at

On current Intel CPUs:

- Registers: 0-1 cycle
- L1 cache: 4 cycles
- L2 cache: 12 cycles
- L3 cache: 26-31 cycles

www.tugraz.at

On current Intel CPUs:

- Registers: 0-1 cycle
- L1 cache: 4 cycles
- L2 cache: 12 cycles
- L3 cache: 26-31 cycles
- DRAM memory: >120 cycles

How every timing attack works:

• learn timing of different corner cases

How every timing attack works:

- learn timing of different corner cases
- later, we recognize these corner cases by timing only

## 1. build two cases: cache hits and cache misses

- 1. build two cases: cache hits and cache misses
- 2. time each case many times (get rid of noise)

- 1. build two cases: cache hits and cache misses
- 2. time each case many times (get rid of noise)
- 3. we have a histogram!

- 1. build two cases: cache hits and cache misses
- 2. time each case many times (get rid of noise)
- 3. we have a histogram!
- 4. find a threshold to distinguish the two cases

1. measure time

- 1. measure time
- 2. access variable (always cache **hit**)

- 1. measure time
- 2. access variable (always cache hit)
- 3. measure time

- 1. measure time
- 2. access variable (always cache hit)
- 3. measure time
- 4. update histogram with delta



1. measure time

- 1. measure time
- 2. access variable (always cache miss)

- 1. measure time
- 2. access variable (always cache miss)
- 3. measure time

- 1. measure time
- 2. access variable (always cache miss)
- 3. measure time
- 4. update histogram with delta

- 1. measure time
- 2. access variable (always cache miss)
- 3. measure time
- 4. update histogram with delta
- 5. flush variable (clflush instruction)

# Time to code

# Accurate timings

- very short timings
- rdtsc instruction: cycle-accurate timestamps

## Accurate timings

- very short timings
- rdtsc instruction: cycle-accurate timestamps

[...] rdtsc function() rdtsc [...]

- do you measure what you *think* you measure?
- **out-of-order** execution → what is really executed

## Accurate timings

- do you measure what you *think* you measure?
- **out-of-order** execution  $\rightarrow$  what is really executed

rdtsc	rdtsc	rdtsc
function()	[]	rdtsc
[]	rdtsc	function()
rdtsc	function()	[]



• use pseudo-serializing instruction rdtscp (recent CPUs)

- use pseudo-serializing instruction rdtscp (recent CPUs)
- and/or use serializing instructions like cpuid

- use pseudo-serializing instruction rdtscp (recent CPUs)
- and/or use serializing instructions like cpuid
- and/or use fences like mfence

- use pseudo-serializing instruction rdtscp (recent CPUs)
- and/or use serializing instructions like cpuid
- and/or use fences like mfence

Intel, How to Benchmark Code Execution Times on Intel IA-32 and IA-64 Instruction Set Architectures White Paper, December 2010.



Michael Schwarz — Security Week Graz 2019

www.tugraz.at





Michael Schwarz — Security Week Graz 2019



• as high as possible

- as high as possible
- most cache hits are below

- as high as possible
- most cache hits are below
- no cache miss below



• Hit  $\rightarrow$  Data is fetched from buffers, L1, L2, or L3

- Hit  $\rightarrow$  Data is fetched from buffers, L1, L2, or L3
- Miss  $\rightarrow$  Data is fetched from DRAM



Michael Schwarz — Security Week Graz 2019

www.tugraz.at



• cache attacks  $\rightarrow$  exploit timing differences of memory accesses
- cache attacks  $\rightarrow$  exploit timing differences of memory accesses
- attacker monitors which lines are accessed, not the content

- cache attacks  $\rightarrow$  exploit timing differences of memory accesses
- attacker monitors which lines are accessed, not the content
- covert channel: two processes communicating with each other

- cache attacks  $\rightarrow$  exploit timing differences of memory accesses
- attacker monitors which lines are accessed, not the content
- covert channel: two processes communicating with each other
  - not allowed to do so, e.g., across VMs

- cache attacks  $\rightarrow$  exploit timing differences of memory accesses
- attacker monitors which lines are accessed, not the content
- covert channel: two processes communicating with each other
  not allowed to do so, e.g., across VMs
- side-channel attack: one malicious process spies on benign processes

- cache attacks  $\rightarrow$  exploit timing differences of memory accesses
- attacker monitors which lines are accessed, not the content
- covert channel: two processes communicating with each other
  not allowed to do so, e.g., across VMs
- side-channel attack: one malicious process spies on benign processes
  - e.g., steals crypto keys, spies on keystrokes

















www.tugraz.at

# Signatures (RSA)

# $M = C^d \mod n$

Michael Schwarz — Security Week Graz 2019

23







Michael Schwarz — Security Week Graz 2019

23

























# Time to code

• locate **key-dependent** memory accesses

- locate **key-dependent** memory accesses
- How to locate key-dependent memory accesses?



• It's complicated:

- It's complicated:
  - Large binaries and libraries (third-party code)

- It's complicated:
  - Large binaries and libraries (third-party code)
  - Many libraries (gedit: 60MB)

- It's complicated:
  - Large binaries and libraries (third-party code)
  - Many libraries (gedit: 60MB)
  - Closed-source / unknown binaries

- It's complicated:
  - Large binaries and libraries (third-party code)
  - Many libraries (gedit: 60MB)
  - Closed-source / unknown binaries
  - Self-compiled binaries

- It's complicated:
  - Large binaries and libraries (third-party code)
  - Many libraries (gedit: 60MB)
  - Closed-source / unknown binaries
  - Self-compiled binaries
- Difficult to find all exploitable addresses

• Preprocessing step to find exploitable addresses automatically

**Exploitation Phase** 

- Preprocessing step to find exploitable addresses automatically
  - w.r.t. "events" (keystrokes, encryptions, ...)

**Exploitation** Phase

- Preprocessing step to find exploitable addresses automatically
  - w.r.t. "events" (keystrokes, encryptions, ...)
  - called "Cache Template"

**Exploitation** Phase

- Preprocessing step to find exploitable addresses automatically
  - w.r.t. "events" (keystrokes, encryptions, ...)
  - called "Cache Template"

**Exploitation Phase** 

• Monitor exploitable addresses

#### Attacker address space





#### Victim address space



#### Cache is empty



Attacker triggers an event
## **Profiling Phase**



Attacker checks one address for cache hits ("Reload")

## **Profiling Phase**



### Update number of cache hits per event

## **Profiling Phase**



### Attacker flushes shared memory



### Repeat for higher accuracy



### Continue with next address



### Continue with next address

Michael Schwarz — Security Week Graz 2019

```
www.tugraz.at 📕
```

```
$> ps -A | grep gedit
$> cat /proc/<pid>/maps
00400000-00489000 r-xp 00000000 fd:01 396356
/usr/bin/gedit
7f5a96991000-7f5a96a51000 r-xp 00000000 fd:01 399365
/usr/lib/x86_64-linux-gnu/libgdk-3.so.0.2200.30
...
```

memory range, access rights, offset, -, -, file name

```
$> cd practicals/02_cache_template_attacks/
$> make
$> # start the targeted program (e.g., gedit)
$> sleep 2; ./profiling /usr/lib/x86_64-linux-gnu/
libgdk-3.so.0.2200.30
```

... and hold down a key in the target program

www.tugraz.at

```
$> cd practicals/02_cache_template_attacks/
$> make
$> # start the targeted program (e.g., gedit)
$> sleep 2; ./profiling /usr/lib/x86_64-linux-gnu/
libgdk-3.so.0.2200.30
```

... and hold down a key in the target program save addresses with peaks!

www.tugraz.at



<pre>\$&gt; # ./spy <file> <offset></offset></file></pre>
<pre>\$&gt; ./spy /usr/lib/x86_64-linux-gnu/libgdk-3.so.0.2200.30 336896</pre>
Monitoring offset 336896
Hit #O
Hit #1
Hit #2

# Time to code

<b>F</b> _	Terminal		- 0	×	Open 🗸	+	Untitled	Document 1	Save	- 1	+	×
File Edit View Search Terminal Help												
% sleep 2; ./spy 300 7f0514 8050 //	40a4000-7f051417b000 usr/lib/x86_64-linux·	r-xp 0x20000 08: gnu/gedit/libged	02 20 lit.s	6 0	1							
() (conce)		-conte 00.01.2017 1- -eOIR->14.03.2017 21	1-44-2	26.2								
File Edit View Search Terminal Help shark‰ ./spy []												
//nome/gamei/iə:							Plain Text 👻	Tab Width: 2 🔻	Ln 1, Col 1		11	NS

#### Cache Template Attack Demo

## **Profiling Phase: 1 Event, 1 Address**

ADDRESS 00202020



## **Profiling Phase: 1 Event, 1 Address**





### Example: Cache Hit Ratio for (0x7c800, n): 200 / 200

### **Profiling Phase: All Events, 1 Address**



## **Profiling Phase: All Events, 1 Address**



Example: Cache Hit Ratio for (0x7c800, u): 13 / 200

## **Profiling Phase: All Events, 1 Address**



Distinguish n from other keys by monitoring 0x7c800

## **Profiling Phase: All Events, All Addresses**



Michael Schwarz — Security Week Graz 2019

D. Gruss, R. Spreitzer, and S. Mangard. Cache Template Attacks: Automating Attacks on Inclusive Last-Level Caches. In: USENIX Security Symposium. 2015.