

# Examining Control Flow Leakage Attacks on TEEs

Luca Wilke\*, Florian Sieck\* and Thomas Eisenbarth (\*equal contribution) Intel Product Assurance and Security - Tech Sharing - November 5, 2024

- 1. Attacks on TDX
- 2. Finding Control Flow Leakages
- 3. Countermeasures

# Goal: Remove cloud provider from TCB



- AMD SEV-SNP
- Intel TDX
- (ARM CCA)



# Single-Stepping Attacks





# Single-Stepping: The bane of TEEs





Interrupt Latency Attacks



Amplifier

Instruction Counting Attacks



Zero-Stepping Attacks

# History of Single-Stepping



# History of Single-Stepping



# History of Single-Stepping



# Countermeasure in TDX



Hypervisor













# Single-Stepping TDX









• SGX-Step: Check if "Accessed"-bit is set

• SGX-Step: Check if "Accessed"-bit is set  $\Rightarrow$  2nd level page tables inaccessible with TDX

- SGX-Step: Check if "Accessed"-bit is set  $\Rightarrow$  2nd level page tables inaccessible with TDX
- SEV-Step: Performance monitoring counters

- SGX-Step: Check if "Accessed"-bit is set  $\Rightarrow$  2nd level page tables inaccessible with TDX
- SEV-Step: Performance monitoring counters  $\Rightarrow$  isolated by TDX

- SGX-Step: Check if "Accessed"-bit is set  $\Rightarrow$  2nd level page tables inaccessible with TDX
- SEV-Step: Performance monitoring counters  $\Rightarrow$  isolated by TDX
- Cache attack on TD code page: works

# Interlude: Zero-Stepping Attacks

## Timer-based Zero-Stepping Attacks



#### Idea

- 1. Repeatedly trigger context switch without state change
- 2. Leak information from context switch
- 3. Example: RAPL software power measurements

## Timer-based Zero-Stepping Attacks



#### Idea

- 1. Repeatedly trigger context switch without state change
- 2. Leak information from context switch
- 3. Example: RAPL software power measurements

## Timer-based Zero-Stepping Attacks



### Idea

- 1. Repeatedly trigger context switch without state change
- 2. Leak information from context switch
- 3. Example: RAPL software power measurements



# StumbleStepping TDX

# StumbleStepping



# StumbleStepping



# StumbleStepping



12/40
# StumbleStepping



- 1. StumbleStepping needs high frequency cache observations
- 2. Regular Flush+Reload blocked by MKTME
- 3. But also: MKTME coherency mechanism enables KeyID-based Flush+Reload

- 1. StumbleStepping needs high frequency cache observations
- 2. Regular Flush+Reload blocked by MKTME
- 3. But also: MKTME coherency mechanism enables KeyID-based Flush+Reload

- 1. StumbleStepping needs high frequency cache observations
- 2. Regular Flush+Reload blocked by MKTME
- 3. But also: MKTME coherency mechanism enables KeyID-based Flush+Reload





#### Exploiting MKTME's coherency mechanism



## Exploiting MKTME's coherency mechanism



### Exploiting MKTME's coherency mechanism



 $\cdot$  Cache attack  $\checkmark$ 

- $\cdot\,$  Cache attack  $\checkmark\,$ 
  - Frequency throttling improves temporal resolution even more

- $\cdot\,$  Cache attack  $\checkmark\,$ 
  - Frequency throttling improves temporal resolution even more
- Page faults for precise termination of prevention mode

# **Primitive Evaluation**

#### Secret Dependent Control-Flow











## Exploitation with StumbleStepping



Goals: Filter all zero-steps & show absence of multi-steps



Evaluation target

- 3 configs: 1, 9 and 10 loop iterations
  - corresponds to 8, 56 and 62 instructions
  - $\cdot$  10 000 measurements for each config
- Found no errors
- only 0.8% zero-step events

## Synthetic StumbleStepping Evaluation

#### Goals: Evaluate Accuracy of inferred instruction count



Inferred Executed Instructions

## Synthetic StumbleStepping Evaluation

#### Goals: Evaluate Accuracy of inferred instruction count



Inferred Executed Instructions

However...

# Synthetic StumbleStepping Evaluation cont.



#### Noise grows with observation length

#### StumbleStepping E2E Attack Sneak Peek



Control Flow events for secp160r1 ECDSA in wolfSSL

Finding Control Flow Leakages

#### Finding Control-Flow Leakages with Microwalk



Microwalk

1. Goal: Get random nonce k < n

- 1. Goal: Get random nonce k < n
- 2. Also: Get it fast

- 1. Goal: Get random nonce k < n
- 2. Also: Get it fast
- 3. Modular reduction approach:

- 1. Goal: Get random nonce k < n
- 2. Also: Get it fast
- 3. Modular reduction approach:
  - 3.1 Sample candidate nonce k'

- 1. Goal: Get random nonce *k* < *n*
- 2. Also: Get it fast
- 3. Modular reduction approach:
  - 3.1 Sample candidate nonce k'
  - 3.2 Compute k as  $k' \mod n$











# Attack Case Study
#### ECDSA Nonce Truncation in wolfSSL

```
int _sp_div_impl(sp_int* a, d, r, trial) {
 1
2
        for (i = a->used - 1; i >= d->used; i--) {
3
            //Calculate trial quotient
            t = sp div word(a -> dp[i], a -> dp[i-1], dt);
4
            do {
5
                 for (j = 0; j < d->used; j++) {...}
6
7
                 for (i = d->used: i > 0: i--)
                     //Event W<sub>2</sub>
8
9
                     if (trial->dp[i] != a->dp[i + o])
10
                          break:
                 if (trial -> dp[i] > a -> dp[i + o]) \{ t --: \}
11
12
                 //Event W_1
            } while (trial->dp[i] > a->dp[i + o]):
13
        }
14
15
   }:
```





27/40



27/40



27/40

Nonce bit distribution given leaked loop iterations for *secp160r1* and brainpoolP224r1.



## Leakage Overview in wolfSSL and OpenSSL

	wolfSSL		OpenSSL	
Curve	Event	MI / FB	Event	MI / FB
	$(W_1, W_2)$ Pr[A = a]	[bit / bit]	$\begin{array}{l} (O_1, O_2) \\ Pr[A = a] \end{array}$	[bit / bit]
bp224r1	(2, *) 0.09	1.6 / 1	(1,0) 1.6 · 10 <sup>-4</sup>	7 / 6
bp320r1	(3, *) < 0.002	3/3	(2, *) 1.7 · 10 <sup>-3</sup>	3/3
bp384r1	(2, *) 0.05	3.5 / 0	(1, *) 0.05	3.5 / 0
secp160r1	(2, *) 1.5 · 10 <sup>-5</sup>	15.6 / 15	(1, *) 1.3 · 10 <sup>-5</sup>	15.8 / 15

#### StumbleStepping the Nonce Bias



#### StumbleStepping the Nonce Bias



Countermeasures

- Only rely on instruction pointer progress
- AEX-Notify shows that reliable "n-stepping" is not possible

- Only rely on instruction pointer progress
- AEX-Notify shows that reliable "n-stepping" is not possible

#### Intel's TDX Module Patch

```
if ((rip_delta > INTEL64_MAX_INST_LEN * 2) || (vcpu_tsc_delta(ld_p) > STEPPING_TSC_THRESHOLD))
// Always use instruction count heuristic if Perfmon is disabled, regardless of TDCS.ATTRIBUTES.
ICSSD
```

if (!perfmon\_enabled)

```
uint64 t inst retired = ia32 rdmsr(IA32 FIXED CTR0 MSR ADDR);
   uint64 t rcx delta = ld p-squest rcx on td entry - ld p-syp ctx.tdyps-squest state.gpr state.
    rcx:
    if ((inst_retired > 1) || ((\emptyset == inst_retired) & (rcx_delta > 1)))
        return FILTER OK CONTINUE:
else if ((rip_delta > INTEL64_MAX_INST_LEN * 2) || (vcpu_tsc_delta(ld_p) > STEPPING_TSC_THRESHOLD))
    return FILTER OK CONTINUE:
```

- StumbleStepping attack won't be mitigated by TDX module
- Protecting crypto code via constant time programming is feasible
- Protecting databases, image decoding, etc. is not feasible

- StumbleStepping attack won't be mitigated by TDX module
- Protecting crypto code via constant time programming is feasible
- Protecting databases, image decoding, etc. is not feasible

- StumbleStepping attack won't be mitigated by TDX module
- Protecting crypto code via constant time programming is feasible
- Protecting databases, image decoding, etc. is not feasible

- StumbleStepping attack won't be mitigated by TDX module
- Protecting crypto code via constant time programming is feasible
- Protecting databases, image decoding, etc. is not feasible

Need principled mitigations as part of the TEE

## Approach 1: Changing MTF flag



## Approach 1: Changing MTF flag



Implementation effort hard to judge for us

Recap Single-Stepping:



Recap Single-Stepping:



## Approach 2: AEX-Notify

Recap Single-Stepping:



## Approach 2: AEX-Notify

Recap Single-Stepping:



Without slowing down Instr 1, reliable, repeated single-stepping is not possible

## Approach 2: AEX-Notify cont.

AEX-Notify idea:



## Approach 2: AEX-Notify cont.

AEX-Notify idea:



- Prefetch code is constant time
- Small, atomic part at the end

- 1. Build on existing interrupt injection mechanisms
- 2. Execute AEX-Notify prefetch as part of each interrupt
- 3. For VM enter without interrupt injection: Force injection of dummy interrupt

**Pro**: Minimial HW-RoT changes  $\Rightarrow$  portable across CVMs? **Con**: Overhead of prefetch, Zero-Stepping security?

- 1. Build on existing interrupt injection mechanisms
- 2. Execute AEX-Notify prefetch as part of each interrupt
- 3. For VM enter without interrupt injection: Force injection of dummy interrupt

**Pro**: Minimial HW-RoT changes  $\Rightarrow$  portable across CVMs? **Con**: Overhead of prefetch, Zero-Stepping security?

- 1. Build on existing interrupt injection mechanisms
- 2. Execute AEX-Notify prefetch as part of each interrupt
- 3. For VM enter without interrupt injection: Force injection of dummy interrupt

**Pro**: Minimial HW-RoT changes  $\Rightarrow$  portable across CVMs? **Con**: Overhead of prefetch, Zero-Stepping security?

- 1. Build on existing interrupt injection mechanisms
- 2. Execute AEX-Notify prefetch as part of each interrupt
- 3. For VM enter without interrupt injection: Force injection of dummy interrupt

**Pro**: Minimial HW-RoT changes  $\Rightarrow$  portable across CVMs? **Con**: Overhead of prefetch, Zero-Stepping security?

- 1. Build on existing interrupt injection mechanisms
- 2. Execute AEX-Notify prefetch as part of each interrupt
- 3. For VM enter without interrupt injection: Force injection of dummy interrupt

**Pro**: Minimial HW-RoT changes  $\Rightarrow$  portable across CVMs? **Con**: Overhead of prefetch, Zero-Stepping security?

- Page granular leakage already sufficient to leak image from jpeg decoding
- TDX already protects the 2nd level page tables but exports page blocking API
- Why not completely restrict forcing page faults?

- Page granular leakage already sufficient to leak image from jpeg decoding
- TDX already protects the 2nd level page tables but exports page blocking API
- Why not completely restrict forcing page faults?

- Page granular leakage already sufficient to leak image from jpeg decoding
- TDX already protects the 2nd level page tables but exports page blocking API
- Why not completely restrict forcing page faults?

## Summary

#### Summary

- Attacks on TDX
  - full single-stepping
  - instruction counting via StumbleStepping;
- Finding & Exploiting Control Flow Leakages
  - Microwalk + Distribution Analysis
  - Nonce truncation in wolfSSL and OpenSSL leaks for certain curves
- $\cdot$  Countermeasures
  - Improved MTF flag, AEXNotify for CVMs
  - Preventing page fault side-channel?
- Responsible Disclosure:
  - Intel fixed single-stepping with TDX module 1.5.06 but not StumbleStepping
  - wolfSSL and OpenSSL switched to rejection sampling

# **Backup Slides**