



# DSASSASSIN: Cross-VM Side-Channel Attacks by Exploiting Intel Data Streaming Accelerator

Ben Chen<sup>1</sup>, Kunlin Li<sup>1</sup>, Shuwen Deng<sup>2</sup>, Dongsheng Wang<sup>2</sup>, Yun Chen<sup>1\*</sup>

<sup>1</sup> The Hong Kong University of Science and Technology (Guangzhou)

<sup>2</sup> Tsinghua University

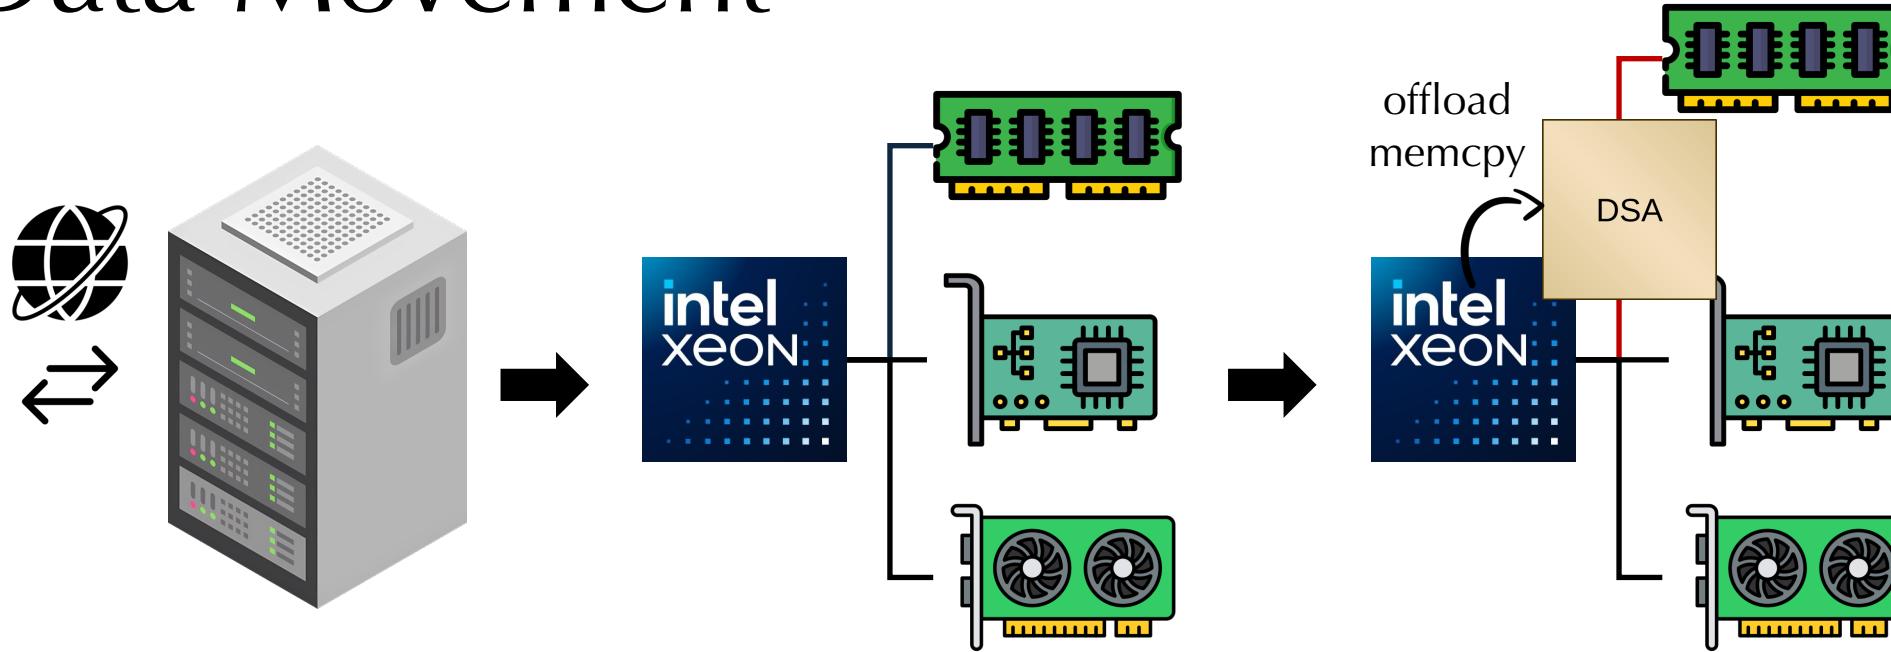


香港科技大学(广州)  
THE HONG KONG  
UNIVERSITY OF SCIENCE AND  
TECHNOLOGY (GUANGZHOU)



清华大学  
Tsinghua University

# Data Movement



Data centers tax for moving data between buffers

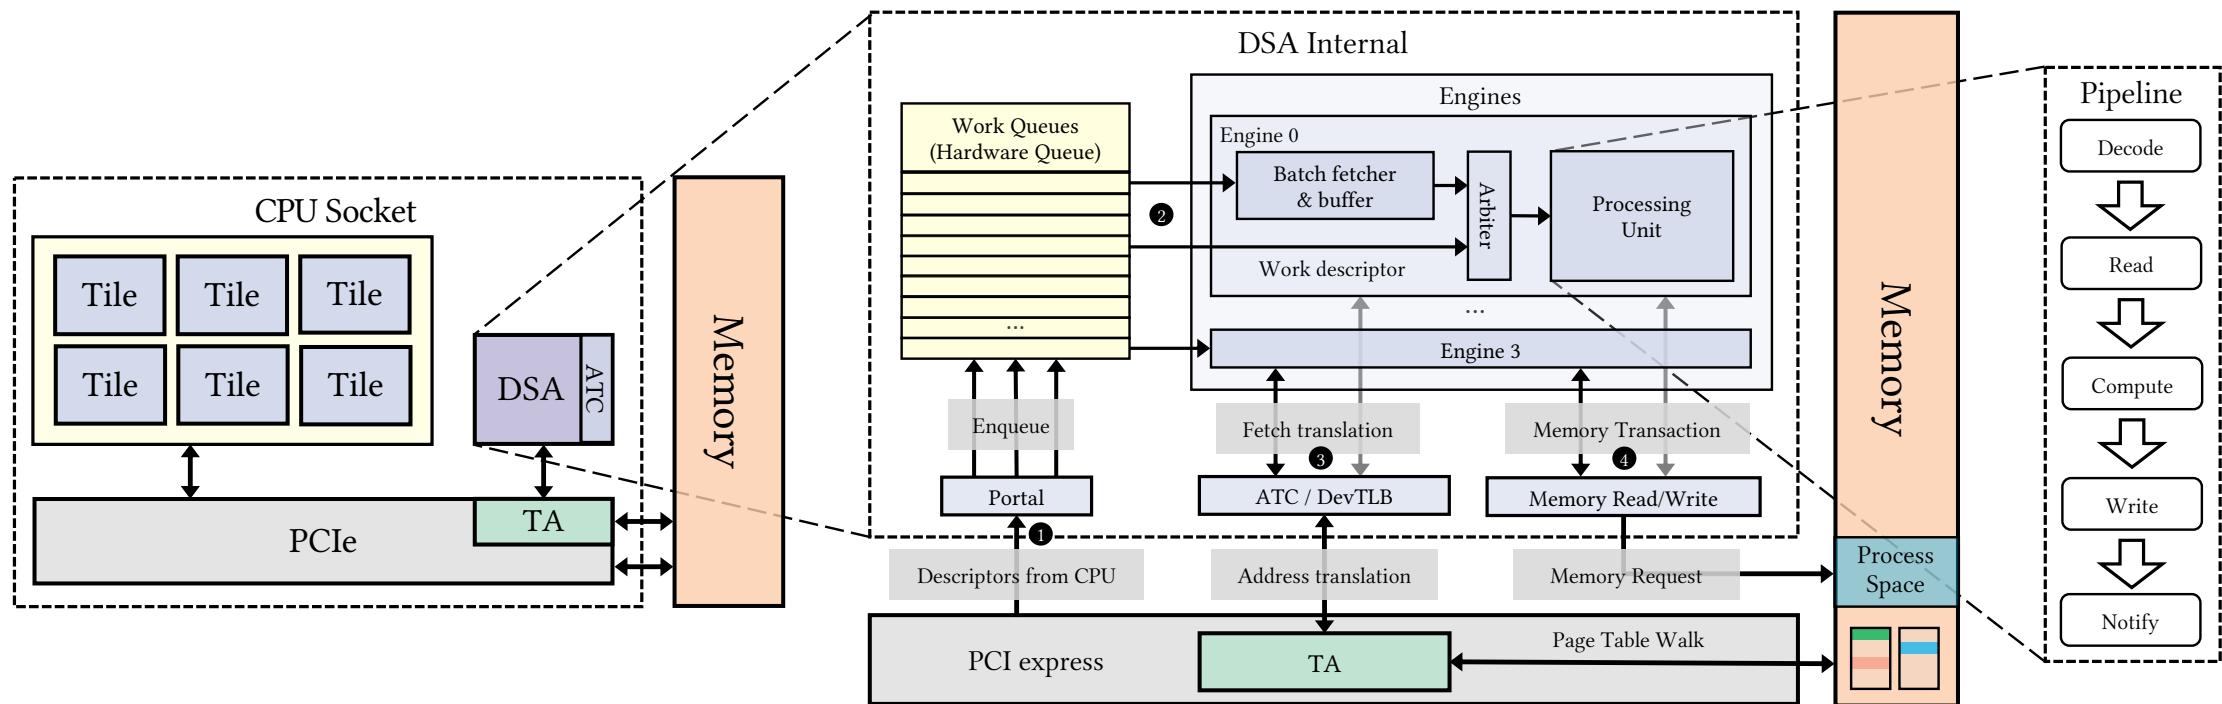
- Dedicated accelerator to save CPU cycles and boost throughput

More devices potentially brings wider attack surfaces

- Example: IOTLB side channels DevIOus [S&P'23] and GPU TLBs [CCS'23]

# Intel Data Streaming Accelerator (DSA)

- Introduced in Intel Xeon 4<sup>th</sup> Gen Scalable and later
- Treated as device w/ Intel VT-d I/O virtualization
- Accelerates data operations such as `memcpy`, `memset`, and `memcmp`
- Introduces asynchronous semantics, saving CPU cycles



# DSA Descriptor

- Descriptor is DSA's instruction
- 64 bytes in total (long enough)
- encodes all necessary information
- does not expose internal states

56	48	40	32	24	16	8	0	Bytes
Opcode	Flags	P	Reserved	PASID	0			
			Completion Record Address					8
			Source Address (src)					16
			Destination Address (dst) / Source2 Address (src2)					24
Reserved	Interrupt Handle		Transfer Size					32
			Destination2 Address (dst2)					40
			Reserved					48
			Unused					56

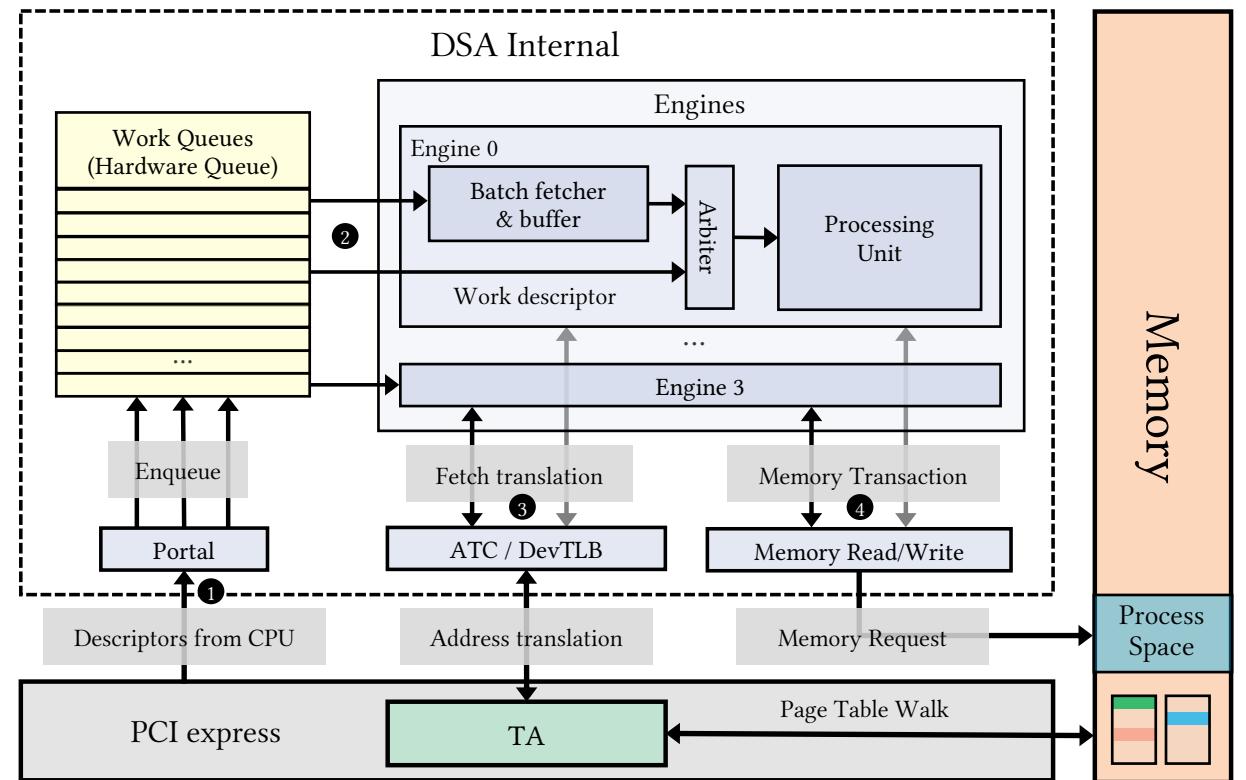
Fig: Descriptor format

56	48	40	32	24	16	8	0	Bytes
Bytes Completed		Unused		Result	Status	0		
			Fault Address					8
			Reserved					16
			Unused					24

Fig: Completion record format

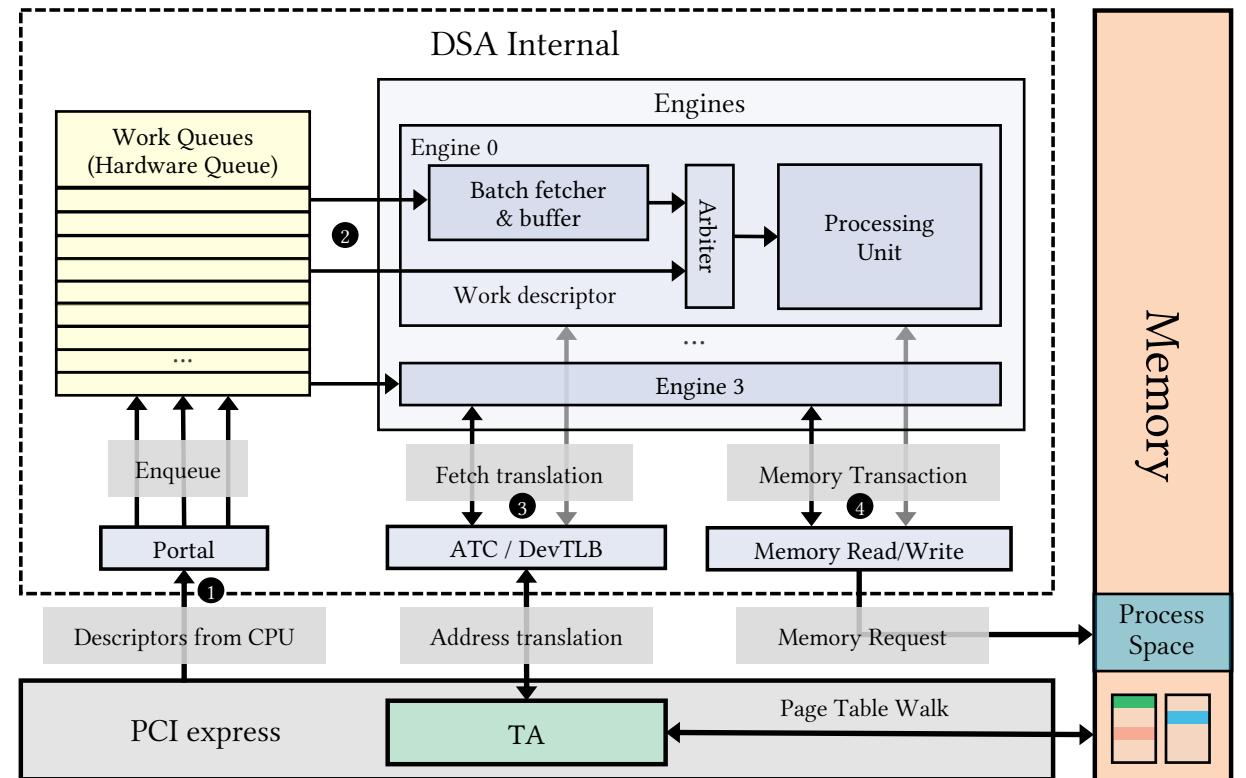
# DSA Workflow

1. User writes descriptor to portal
2. Portal enqueues and dispatches to engine
3. Engine fetches address translation
4. issues memory requests via PA
5. notifies user by writing completion record



# DSA Workflow

- IOMMU => TA (Intel VT-d)
  - Offload translation requests
  - Directly use PA
- I/O VA => Process VA
  - Same VA space
  - Needless to allocate IOVA
  - PASID to walk process's page table

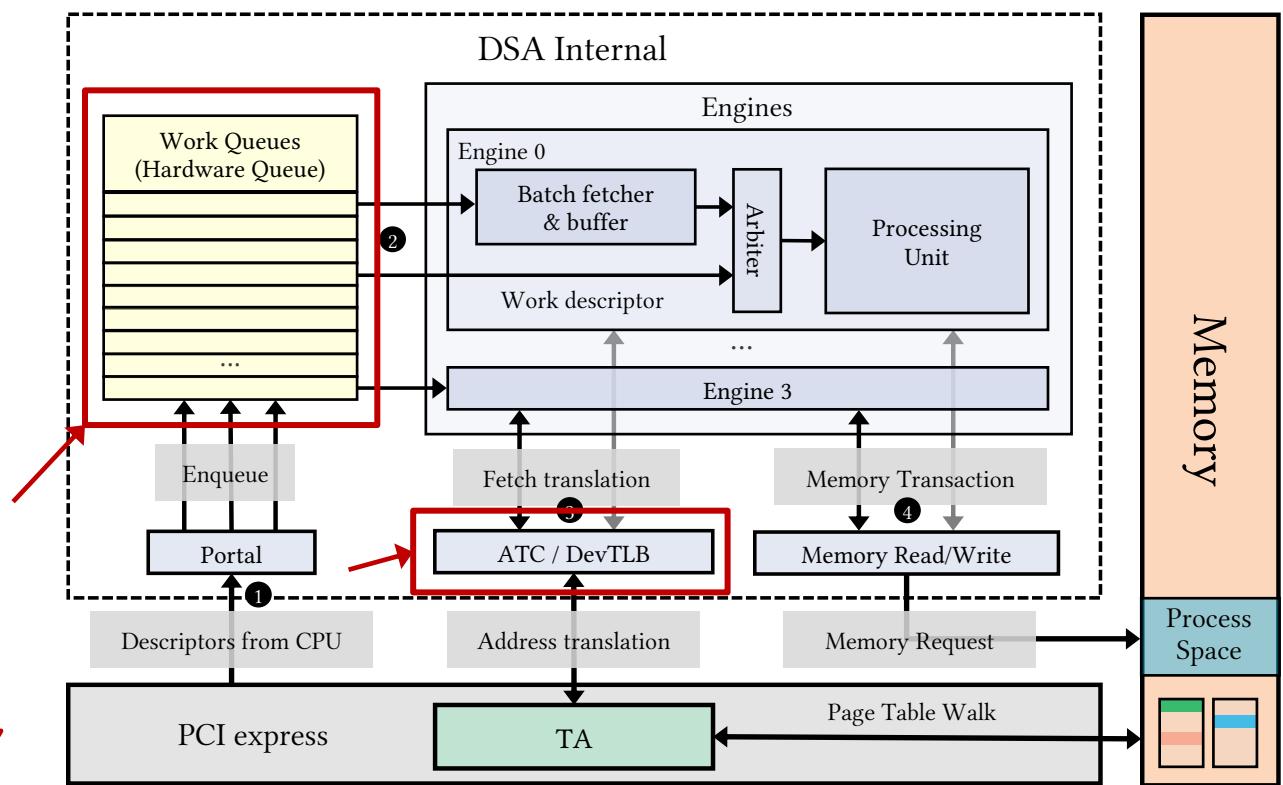


# Micro-arch Side-Channels

Speculative Attacks	<ul style="list-style-type: none"><li>▪ Spectre<ul style="list-style-type: none"><li>▪ PHT, BTB, RSB, etc.</li></ul></li><li>▪ Meltdown<ul style="list-style-type: none"><li>▪ RIDL, ZombieLoad etc.</li></ul></li><li>▪ uarch Optimizations<ul style="list-style-type: none"><li>▪ Prefetcher, predictor (LAP/LVP)</li></ul></li></ul>	<ul style="list-style-type: none"><li>✓ Proactive-triggered</li><li>✓ Low noise</li><li>✗ Single-core</li><li>✗ Need cache or other primitives</li></ul>
	<ul style="list-style-type: none"><li>▪ Caches<ul style="list-style-type: none"><li>▪ LLC, weak coherence</li></ul></li><li>▪ Port contention (e.g. ALU)</li><li>▪ Bus contention</li><li>▪ Buffer (e.g. ROB, scheduler)</li></ul>	<ul style="list-style-type: none"><li>✓ Non-speculative</li><li>✓ Stealthy and hard to patch</li><li>✗ Higher noise</li><li>✗ Single-core (or required SMT co-location)</li></ul>
	<ul style="list-style-type: none"><li>▪ Interconnects<ul style="list-style-type: none"><li>▪ PCIe, NoC, IOTLB etc.</li></ul></li><li>▪ Devices uarch<ul style="list-style-type: none"><li>▪ GPU, NPU, NIC etc.</li><li>▪ <b>This work</b></li></ul></li></ul>	<ul style="list-style-type: none"><li>✓ Cross-VM</li><li>✓ Low noise</li><li>✗ Passive-triggered</li><li>✗ Required high-accurate timer (mostly)</li></ul>

# DSA Is Meant To Be Shared

- Architecturally isolates processes via PASID
  - TA only walks tenant's memory space
  - SWQ synchronizes descriptor submissions
- Enables DSA to share seamlessly
- But, is it safe from side-channels?
- Can we attack its microarchitecture, e.g. DevTLB and SWQ?



# Methodology

- PMU: Perfmon for DSA, supported by Linux perf tool

Event Name	Category	Event No.	Description
EV_ATC_ALLOC	0x2	0x40	Number of requests to DevTLB
EV_ATC_NO_ALLOC	0x2	0x80	Number of requests not allocated DevTLB entry
EV_ATC_HIT_PREV	0x2	0x100	Number of DevTLB hits

- Timer: rdtsc, to benchmark the completion latency
- Primitives used to benchmark DevTLB and SWQ
  - **noop**: write to completion record at **addr**
  - **memcmp**: compare two regions from **src** and **src2**
  - **memcpy**: copy **src** to **dst**
  - **dualcast**: copy **src** to **dst** and **dst2**

# DevTLB Index Policy

```
// base is 4KB aligned
probe_noop(base);
probe_noop(base + OFFSET);
// probe if base is present
probe_noop(base); // miss

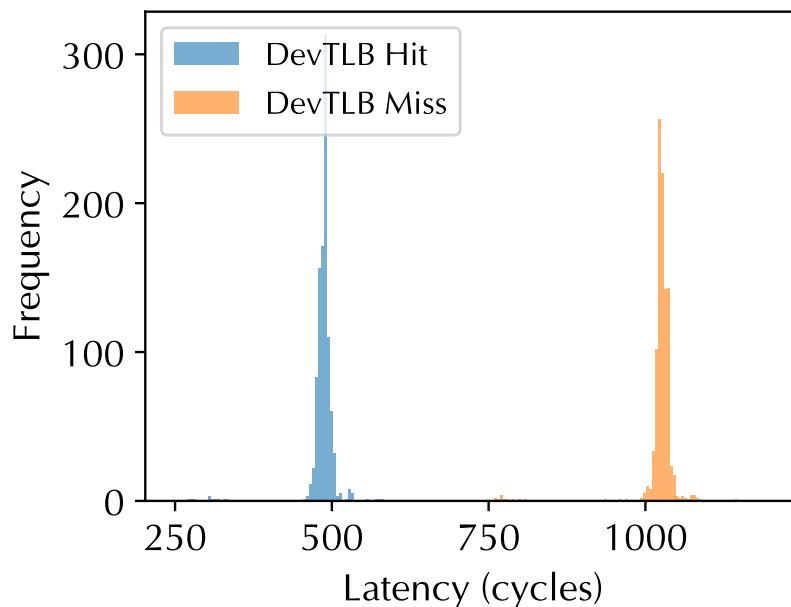
// different pages
assert(PAGE_OF(src0) != PAGE_OF(dst0));
assert(PAGE_OF(src0) != PAGE_OF(src1));
probe_memcpy(src0, dst0);
probe_memcpy(src1, dst0); // 1 hit: src

// src2 == dst
probe_memcmp(src, src2);
probe_memcpy(src, dst); // 1 hit: src
```

- Preliminary: check page boundary
  - Not surprisingly 4KB
- Assume set-associativity
- Immediately evicts base. Weird!
- What if it's not indexed by VA bits?
- Three addresses on different pages
- Entry of dst0 hits, so src is untouched
- Another entry found!
- What if same address is placed at different fields in descriptor?
- Hit! **So DevTLB is indexed by fields**

# Benchmark DevTLB Latency

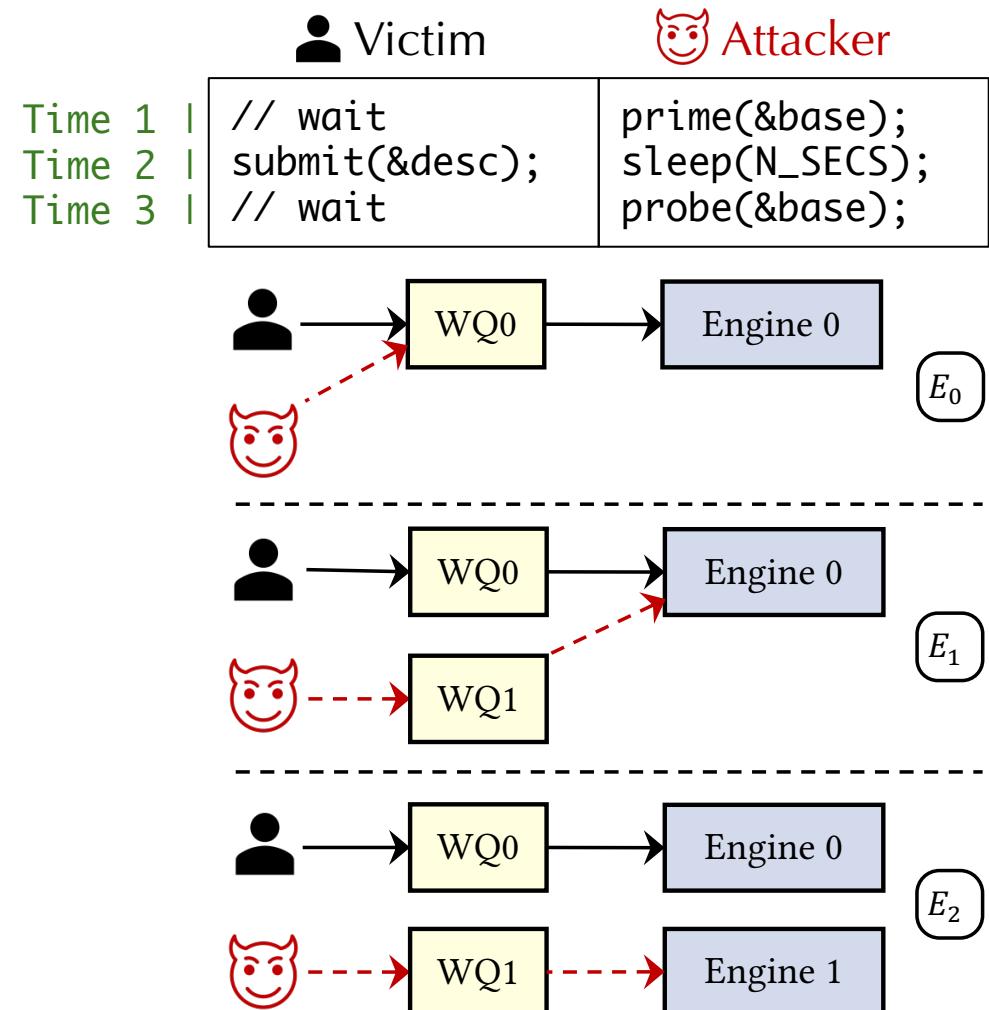
```
void* base = malloc(); base[0] = 0; probe_noop(base); // warm up
for (int i = 0; i < ITERATION; i++) {
    hit[cnt] = time(probe_noop(base)); // record hit latencies
    miss[cnt++] = time(probe_noop(base + OFFSET)); // miss latencies
}
```



- Poll for completion
- DevTLB hit / miss latency is distinguishable with the unprivileged timer `rdtsc`
- ~500 cycles variation
- We can do side-channel

# DevTLB Index & Replacement

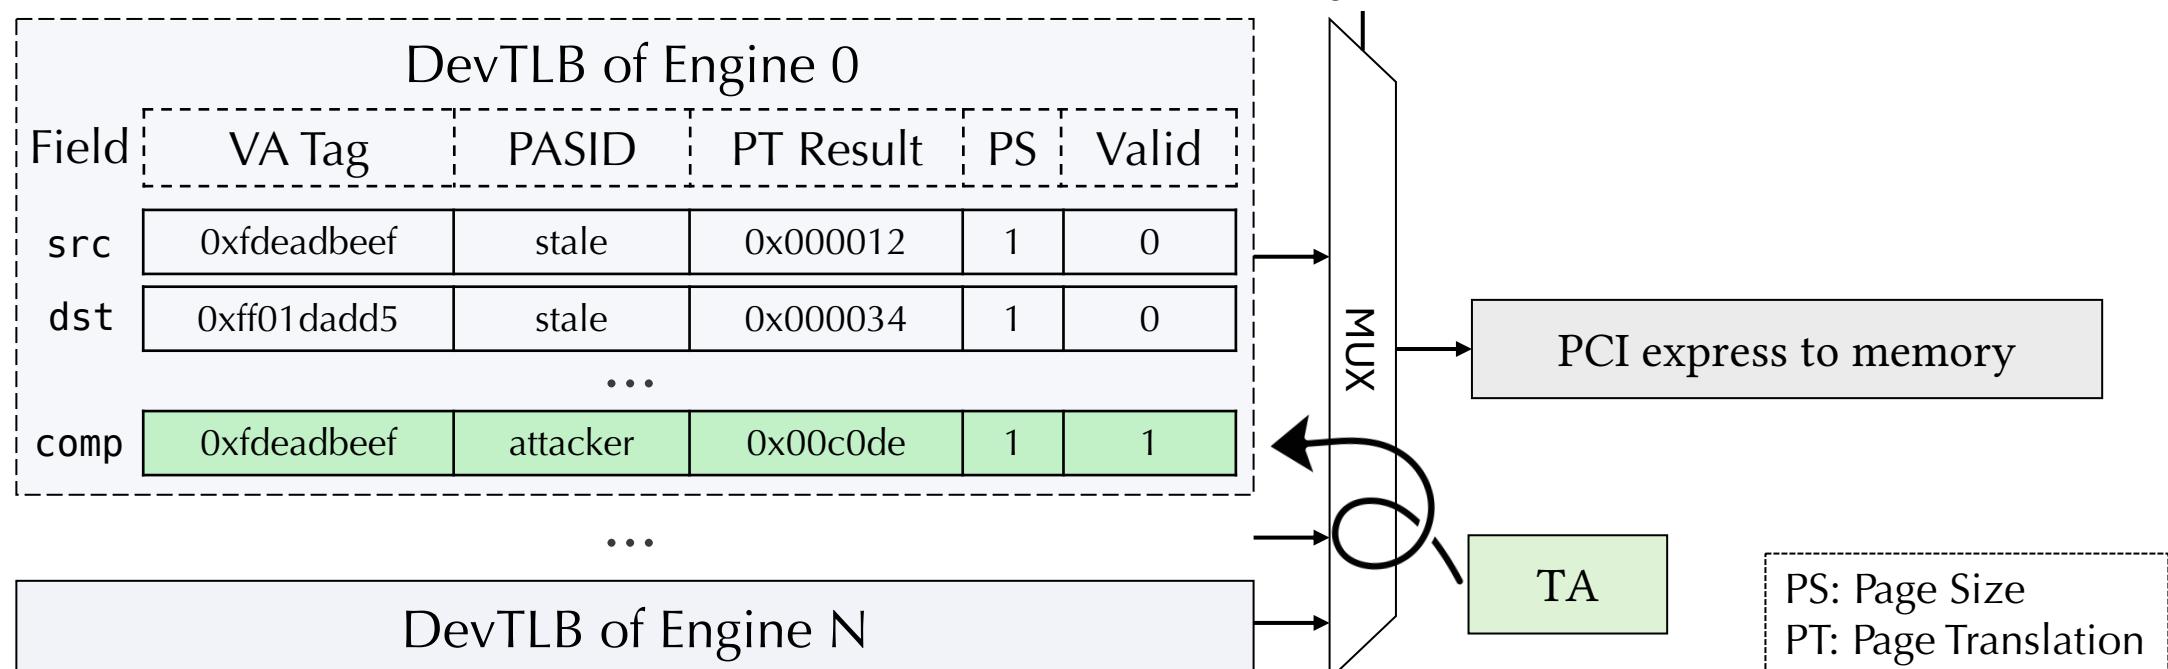
- Two isolated processes:
  - attacker casts Prime+Probe
  - victim submits descriptor during attacker's waiting
- Three configurations
  1. same WQ, same engine
  2. different WQs, same engine
  3. both are separate



# DevTLB Attack

PC  
→ `void* base = malloc(); base[0] = 0;  
probe_noop(base); // attacker  
submit(&wq, &desc); // victim  
time(probe_noop(base)); // attacker`

- First step: attacker primes some entries



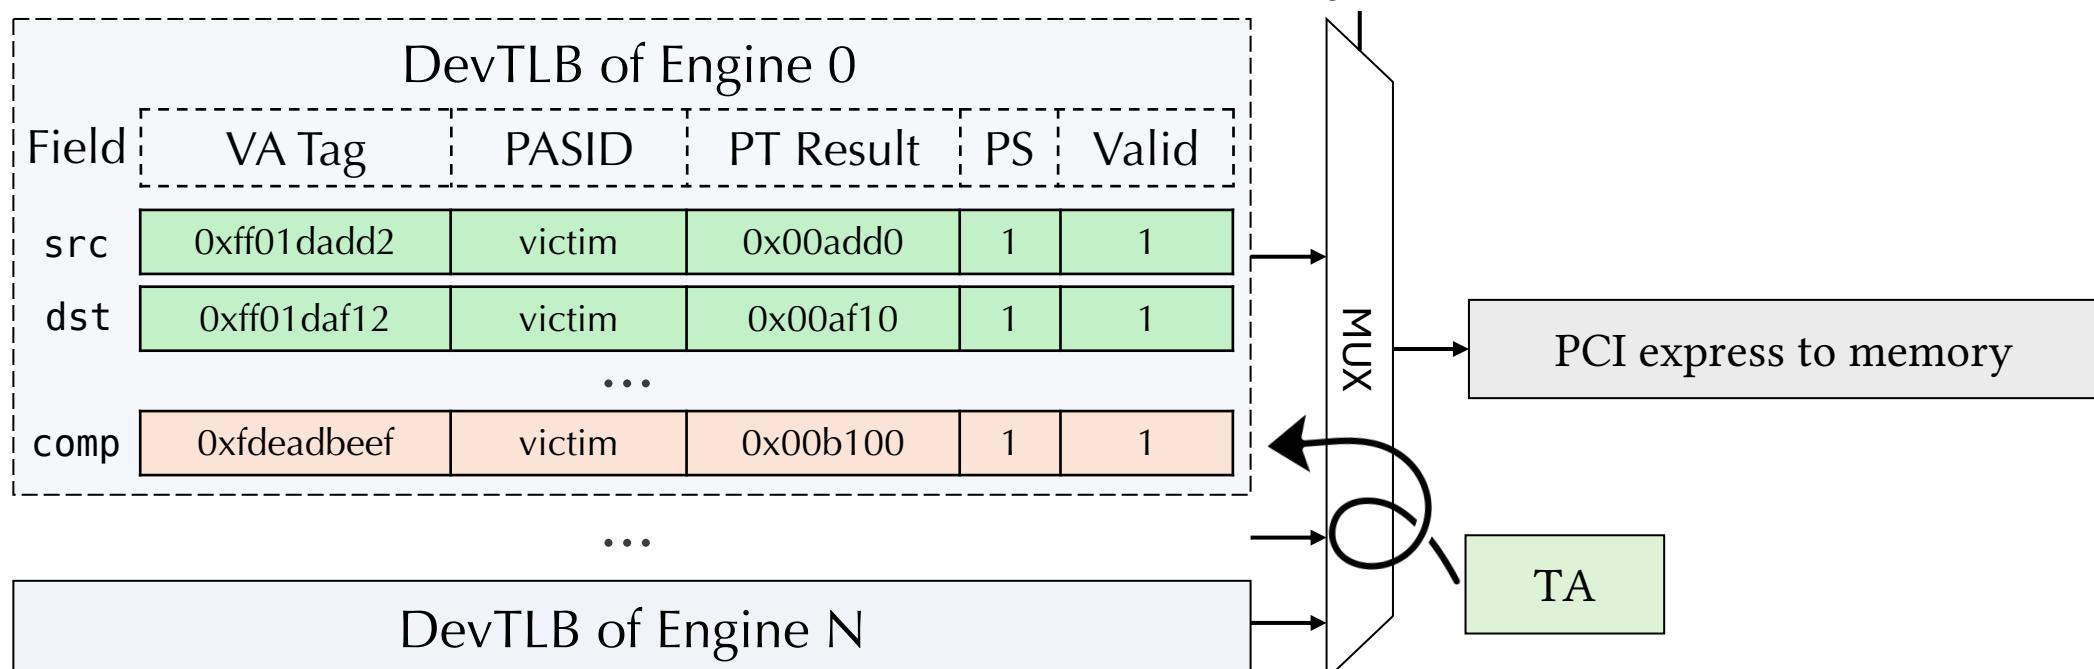
# DevTLB Attack

PC

```
void* base = malloc(); base[0] = 0;
probe_noop(base);           // attacker
submit(&wq, &desc);        // victim
time(probe_noop(base)); // attacker
```

→

- Wait for victim to evict the entries in DevTLB
- No PASID to protect the entries, so  $E_0$  and  $E_1$  succeed

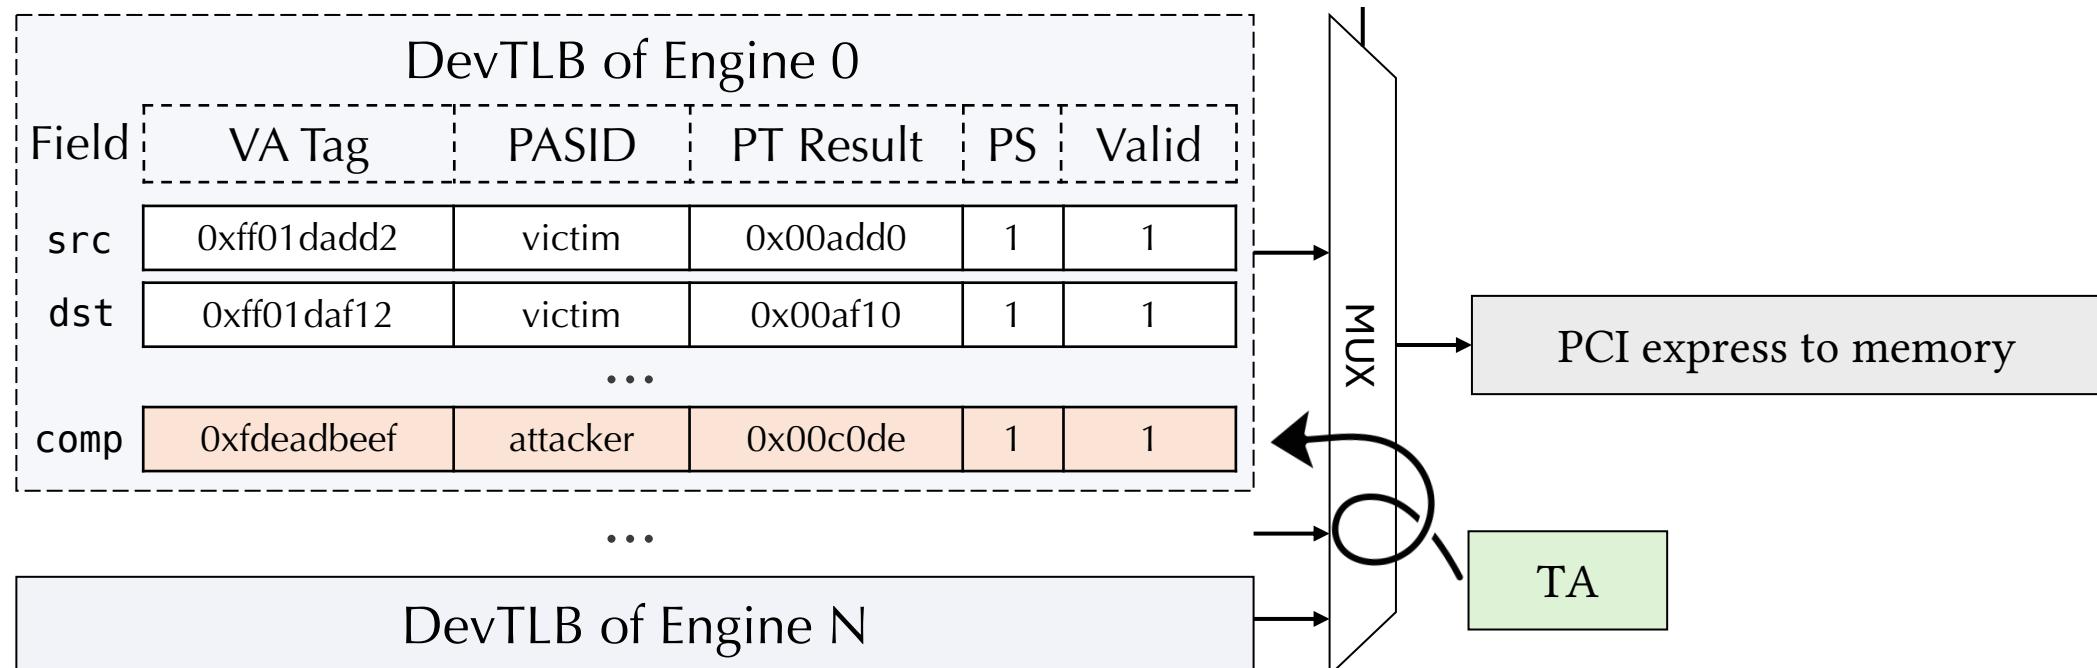


# DevTLB Attack

PC

```
void* base = malloc(); base[0] = 0;
probe_noop(base);           // attacker
submit(&wq, &desc);        // victim
time(probe_noop(base)); // attacker
```

→



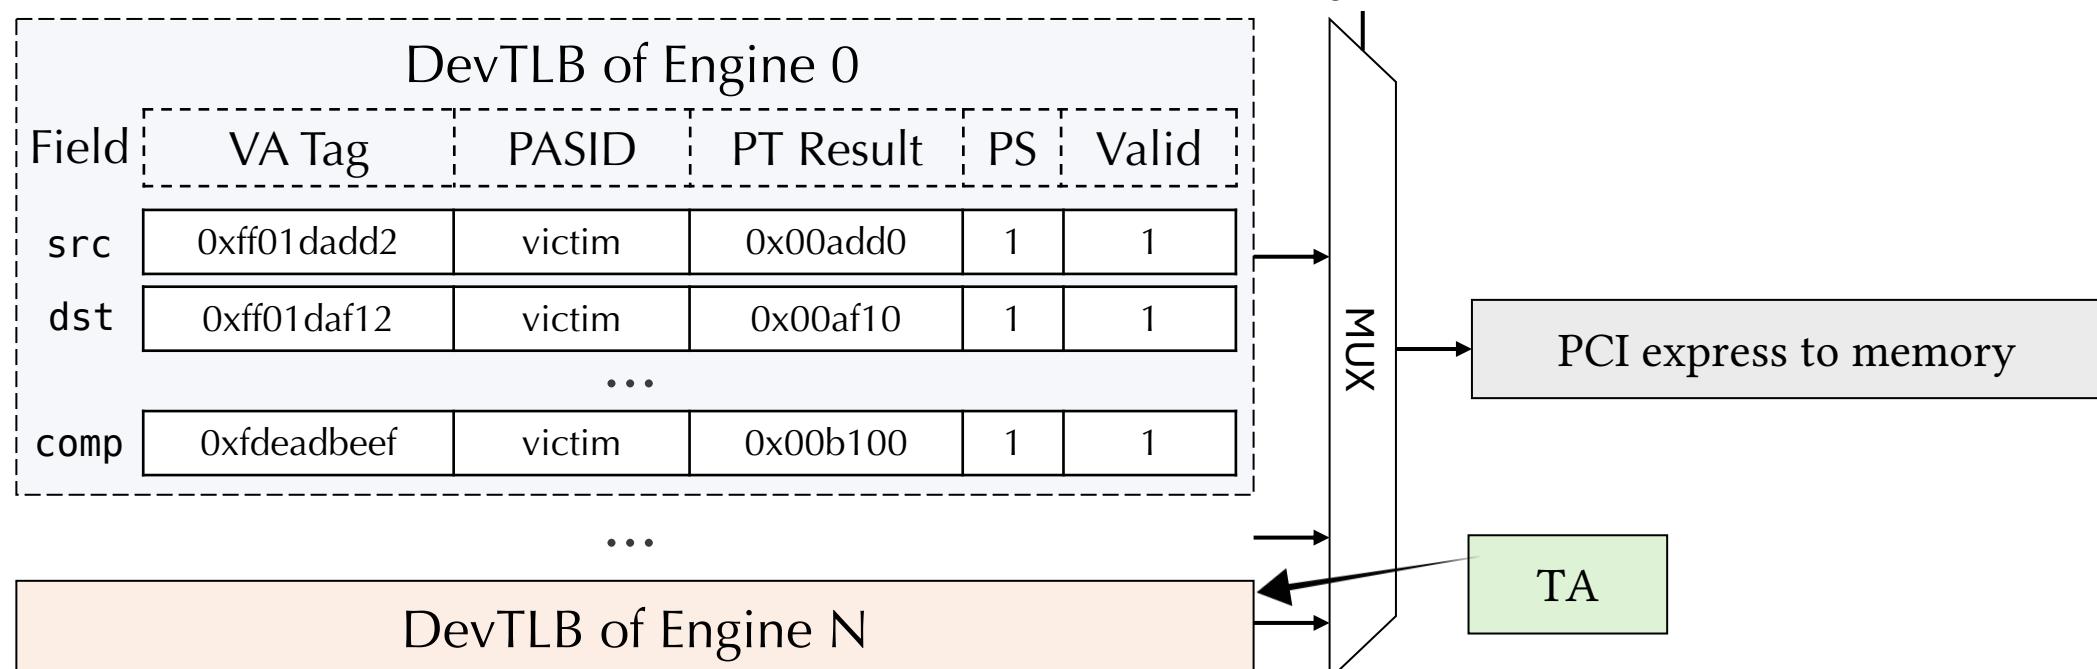
- Time the access
- Longer means the entry was evicted => victim's behavior

# DevTLB Attack

PC

```
void* base = malloc(); base[0] = 0;
probe_noop(base);           // attacker
submit(&wq, &desc);        // victim
time(probe_noop(base)); // attacker
```

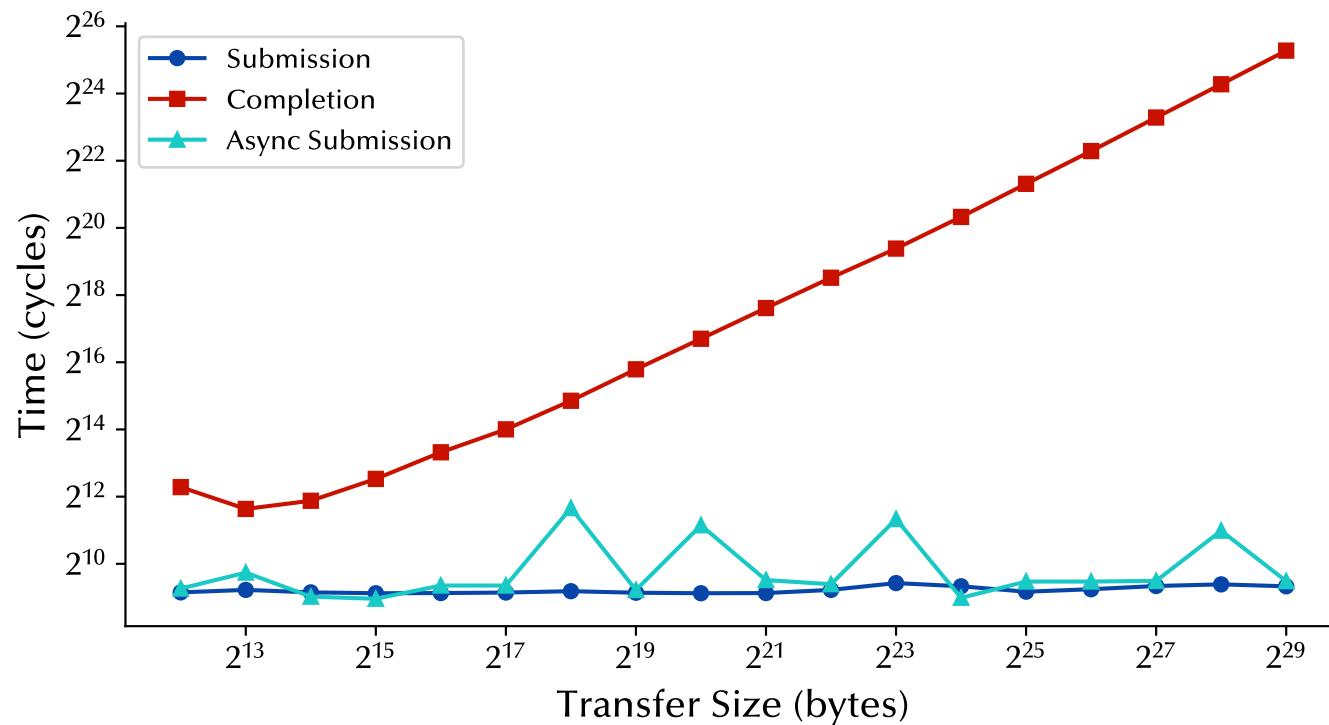
→



- **Engine exclusive:**  $E_2$  failed to evict because engine ID split the entries and the translation was cached in another entry

# Shared Work Queue

- Benchmarking submission (enqueue) and completion latencies
- Completion latency is linear with respect to data transfer size (predictable)
- Submission is constant-time

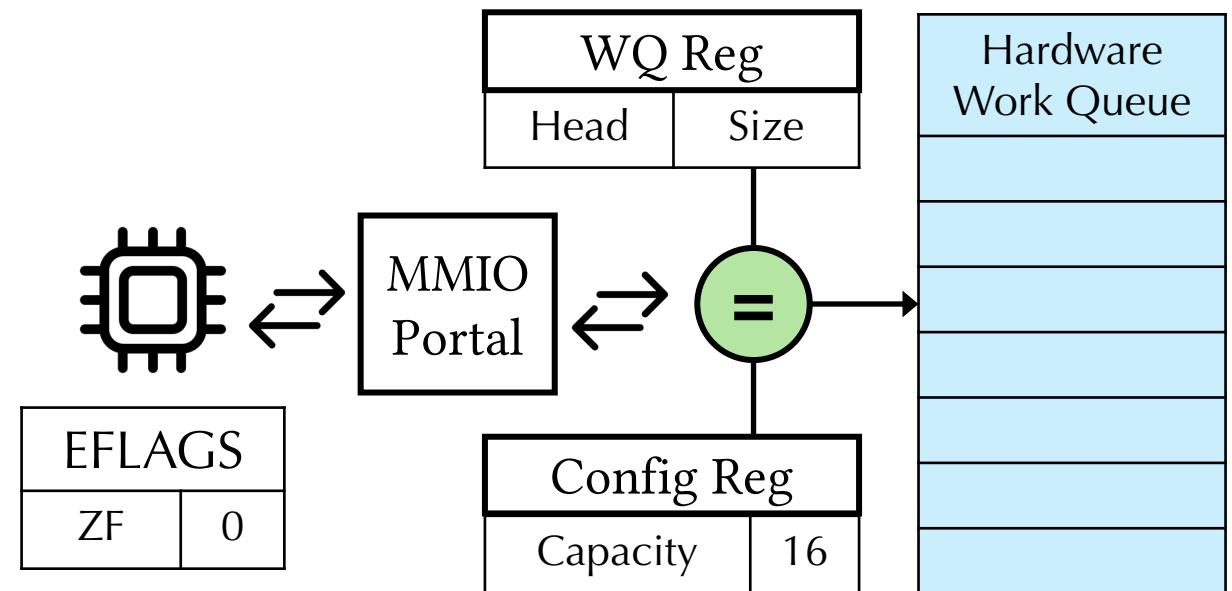


# SWQ Portal

- Deferable Memory Write (DMWr) is non-posted, waiting for feedback
  - Indicates the submission status by setting the bit ZF
  - SWQ checks size with configured capacity
- Architecturally visible SWQ status
  - Enables timer-less side-channel: Congest + Probe

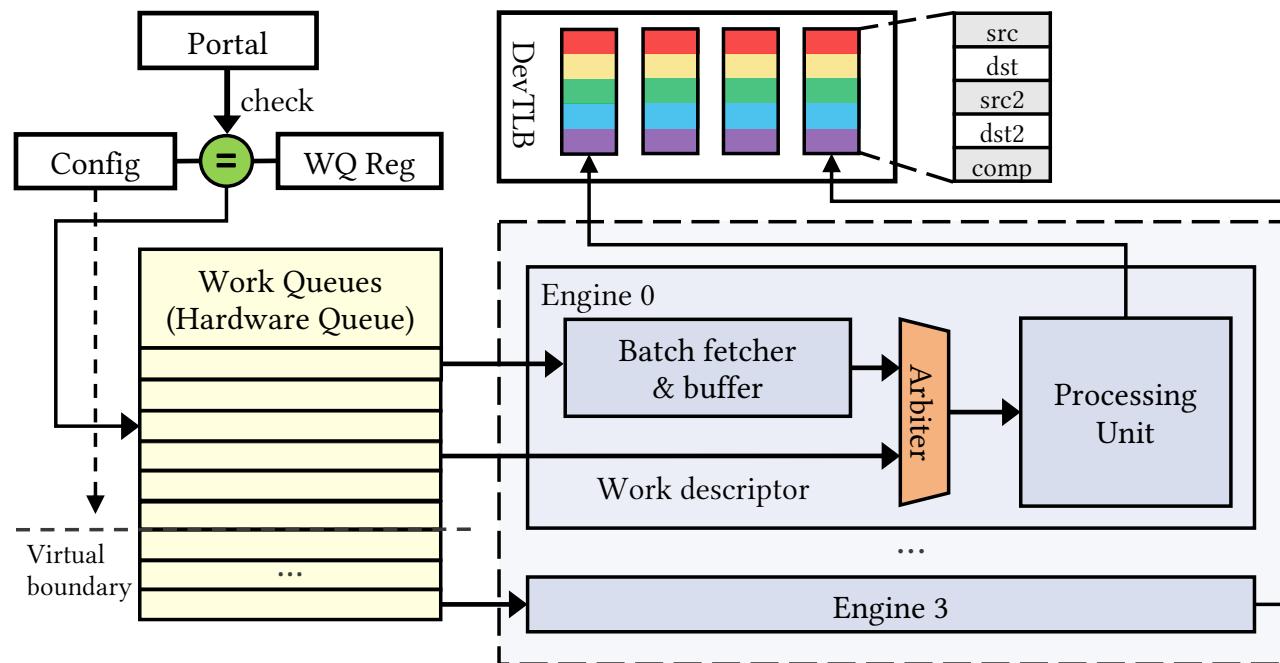
```
void congest() {
    enqcmd(wq, &mass_desc);
    for (i = 0; i < SIZE; i++)
        enqcmd(wq, &desc[i]);
}

bool probe() {
    return enqcmd(wq, &desc);
}
```



# DSA uArch Takeaway

- DevTLB is indexed by field types (5 entries) and engine IDs
- SWQ exposes architectural state and exhibits constant-time behavior
- **Neither of them is PASID-partitioned**



# Attack Primitives

1 Step: Prime (DevTLB) / Congest (SWQ)  
 Attacker

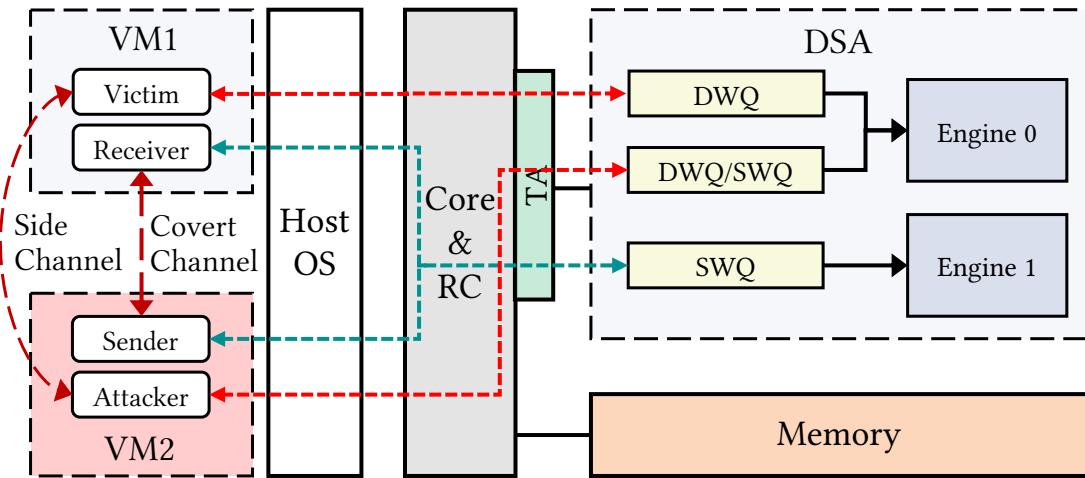
```
probe_noop(base); // prime DevTLB
for (int i = 0; i < WQ_SIZE; i++)
    submit(&swq0, &desc); // congest SWQ
```

2 Step: Leakage via access to DSA  
 Victim

```
// ...
bit = submit(&swq0, &desc);
// ...
```

3 Step: Probe the DevTLB / SWQ  
 Attacker

```
bit = time(probe_noop(base)) > threshold;
bit = submit(&swq0, &desc);
```



- Isolation: attackers and victims run on different VM and cannot cross boundary
- Co-location: mapped to same engine (red path) or same SWQ (teal path)
- Target: DSA applications, e.g. DPDK, DTO (DSA Transparent Offload library)

# DSASSASSIN: Setup & Attacks

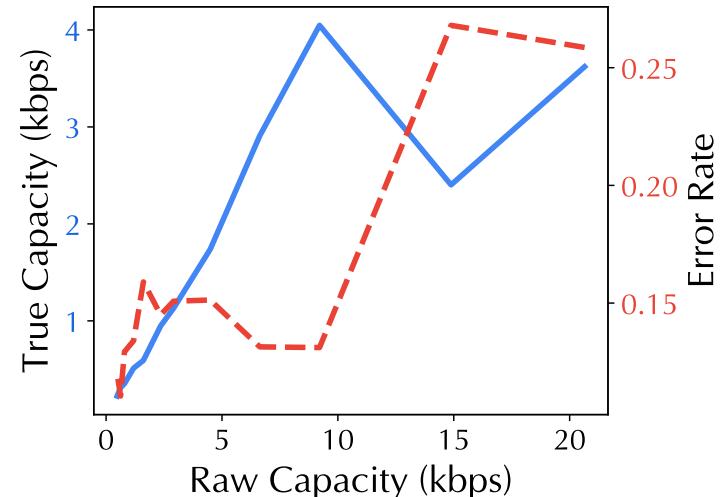
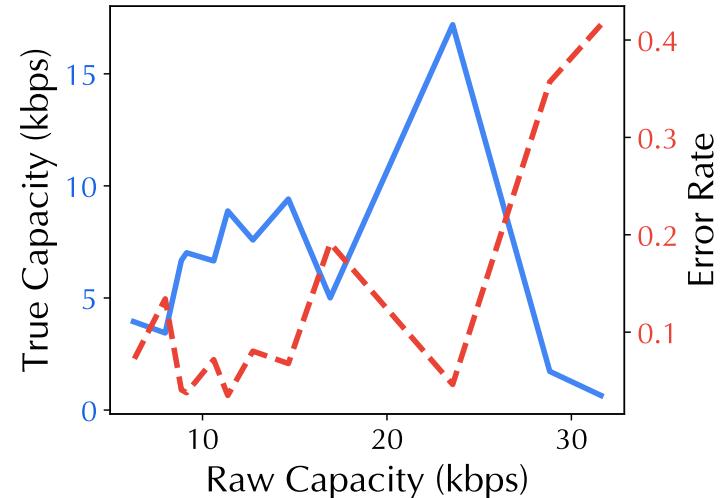
We launched four attacks to showcase the primitives,  $DSA_{DevTLB}$  &  $DSA_{SWQ}$



Specification	Local	Cloud
Provider	Self-hosted	Alibaba Cloud
Processor	Xeon Platinum 8468V	Xeon Platinum 8475B
Architecture	Sapphire Rapids (4 <sup>th</sup> Gen Xeon Scalable)	
OS	Ubuntu 24.04 LTS	
VMM	KVM	
I/O Virtualization	Intel Scalable IOV	Single-Root IOV
DSA Instance	DSA (v1.0) * 2	DSA (v1.0) * 6

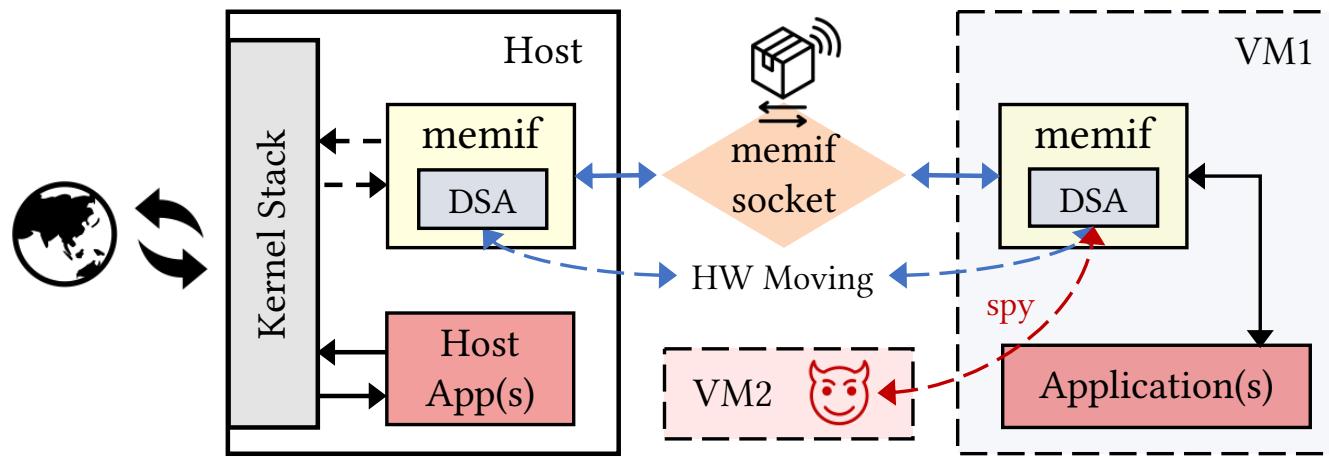
# DSASSASSIN: Covert Channel

- A preliminary test to see how fast we can leak
- Victim deliberately sends message to the attacker under VM isolation
- Protocol:
  - eviction or submission means bit '1'; otherwise '0'
  - synchronization by emitting consecutive bit '1's
- Performance:
  - DevTLB (top fig.): 17.19 Kbps true cap, 4.63% BER
  - SWQ (bottom fig.): 4.02 Kbps true cap, 13.11% BER
  - True capacity larger than IOTLB attacks



# DSASSASSIN: Website Fingerprinting

- User-space network stack VPP upon DPDK
- VPP memory interface (memif) accelerates VM networking
  - Example: Calico VPP on Kubernetes
- memif supports DSA acceleration for memory copying
- We can spy on this



# DSASSASSIN: Website Fingerprinting

- DevTLB miss => packet transmission
- Example: unique network fingerprints when accessing 3 websites
- LSTM model to classify 15 websites on Chrome 138.0: **96.5% accuracy**
- Top 100: **87.6% accuracy**

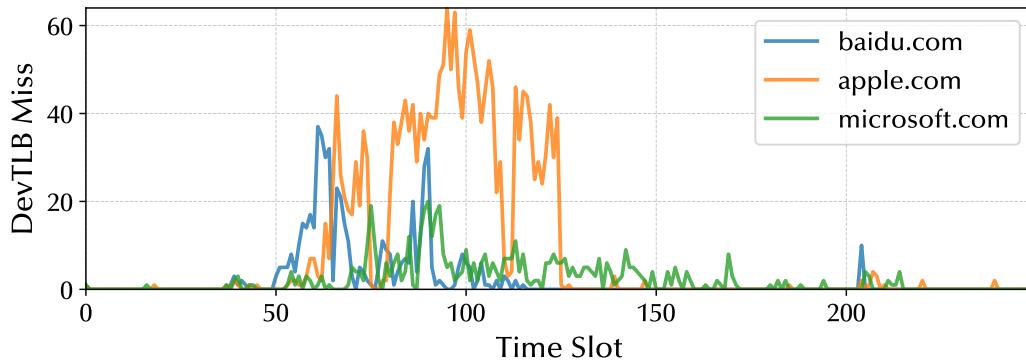
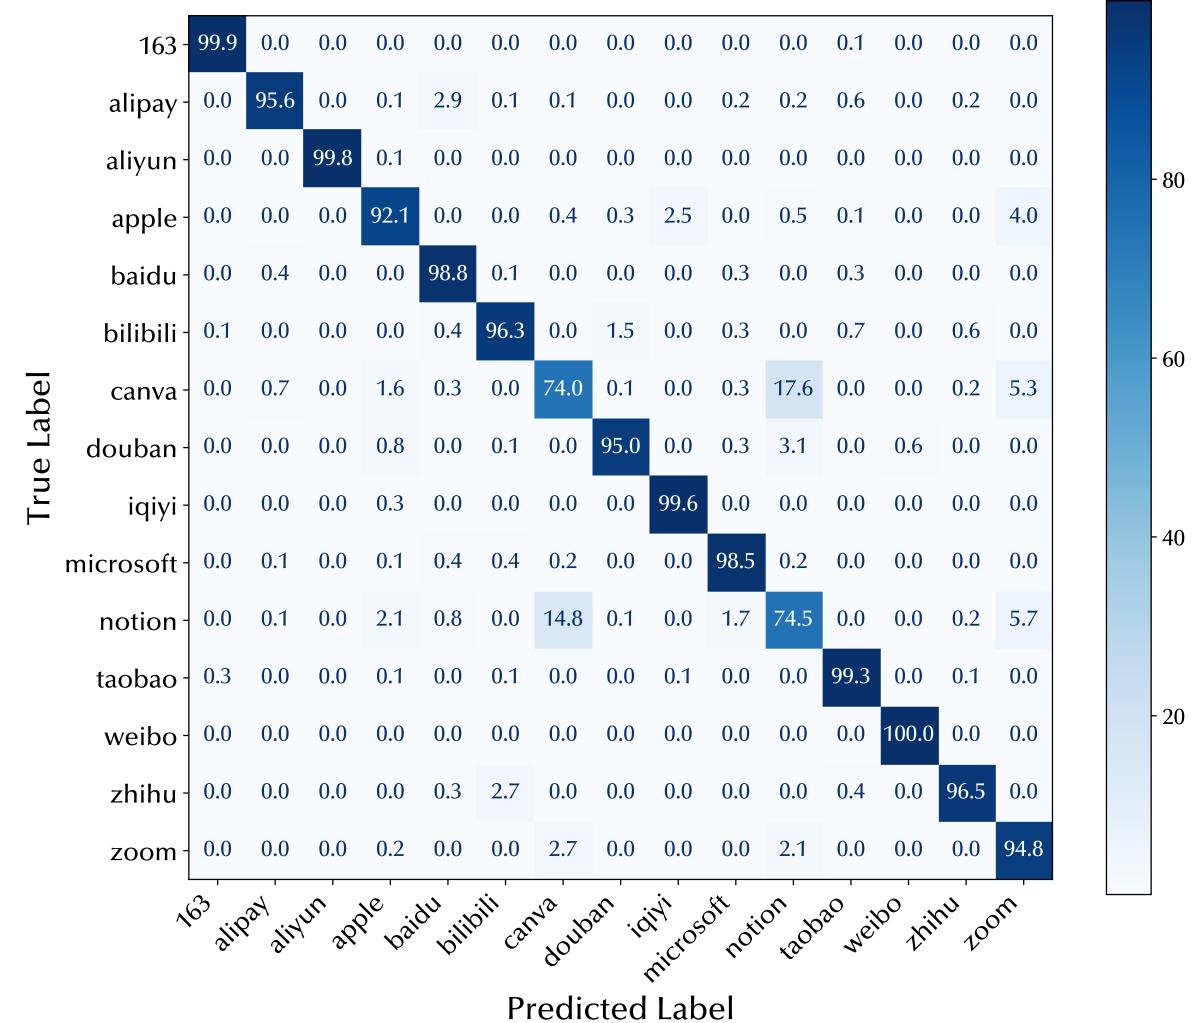
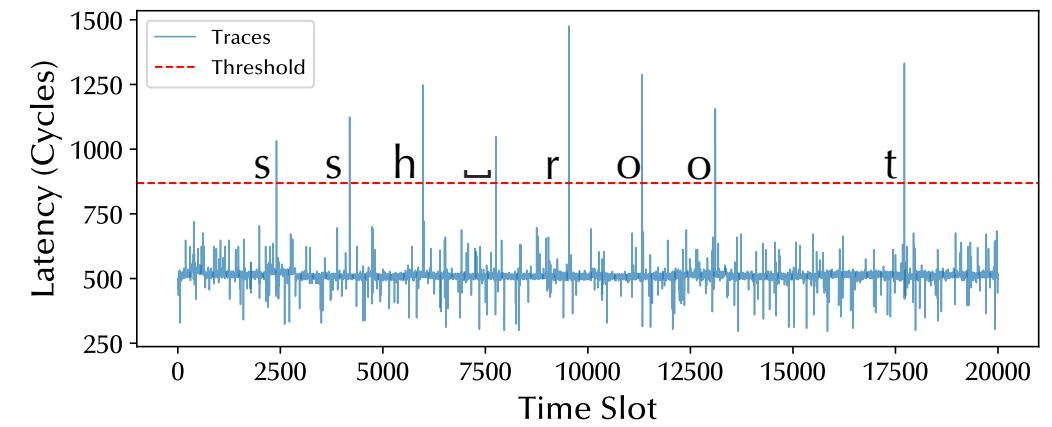
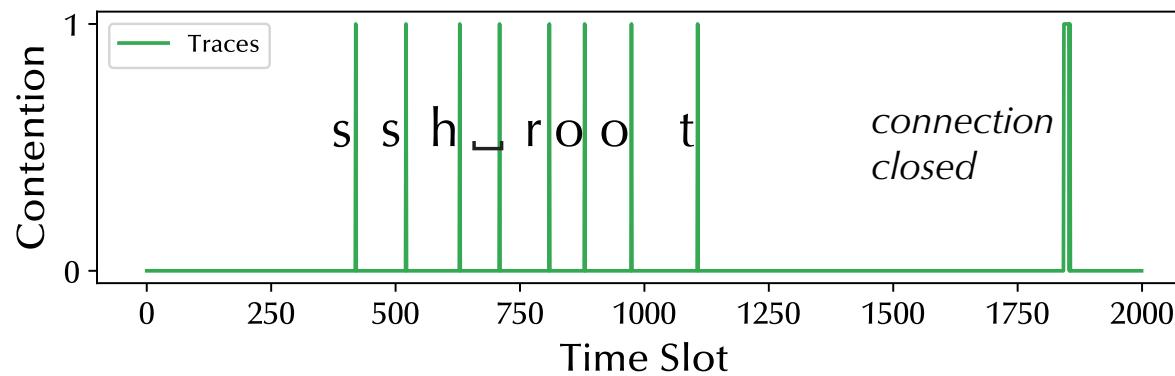


Fig. Example traces



# DSASSASSIN: SSH Keystroke

- Victim's SSH client sends packets upon key pressed
  - buffer copying involved => DTO acceleration
- Attacker can try to learn the exact timing of victim's keystrokes
  - Periodically priming DevTLB or congesting SWQ
  - Shouldn't miss any keystrokes ( $F_1$  score) or deviate too much (std)
  - DevTLB: **92.0%, 5.29 ms**; SWQ: **98.4%, 1.21 ms**

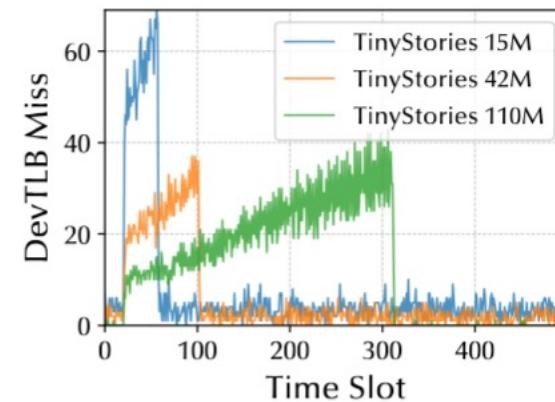


# DSASSASSIN: LLM Fingerprinting

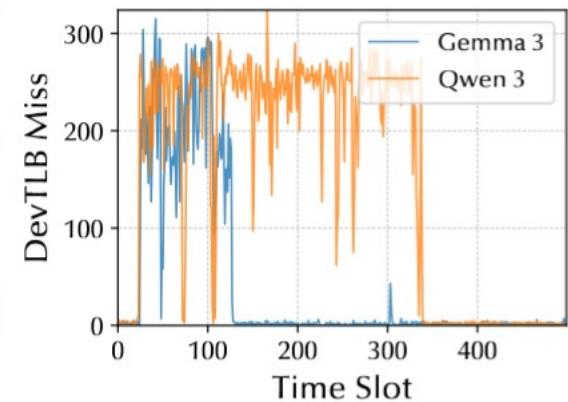
- Lots of data movement in CPU-only and CPU-GPU inference, e.g. weights and activation values
  - May need DSA to accelerate
  - Example: LILo [HPCA'26] applies Intel IAA, DSA's sibling
- Side-channel traces to classify the model's architecture
- We can achieve **98.6% accuracy** on these 8 models

Table. Tested models

Model	Parameters	Description
TinyStories	15M/42M/110M	Tiny model in LLaMA 2
Meta LLaMA 2	7B	Official LLaMA 2
Gemma 3	1B/4B	Model on single GPU
Qwen3	1.7B/4B	Dense and MoE model



(a) TinyStories 15M/42M/110M



(b) Gemma 3 and Qwen 3 (4B)

# Noise Analysis

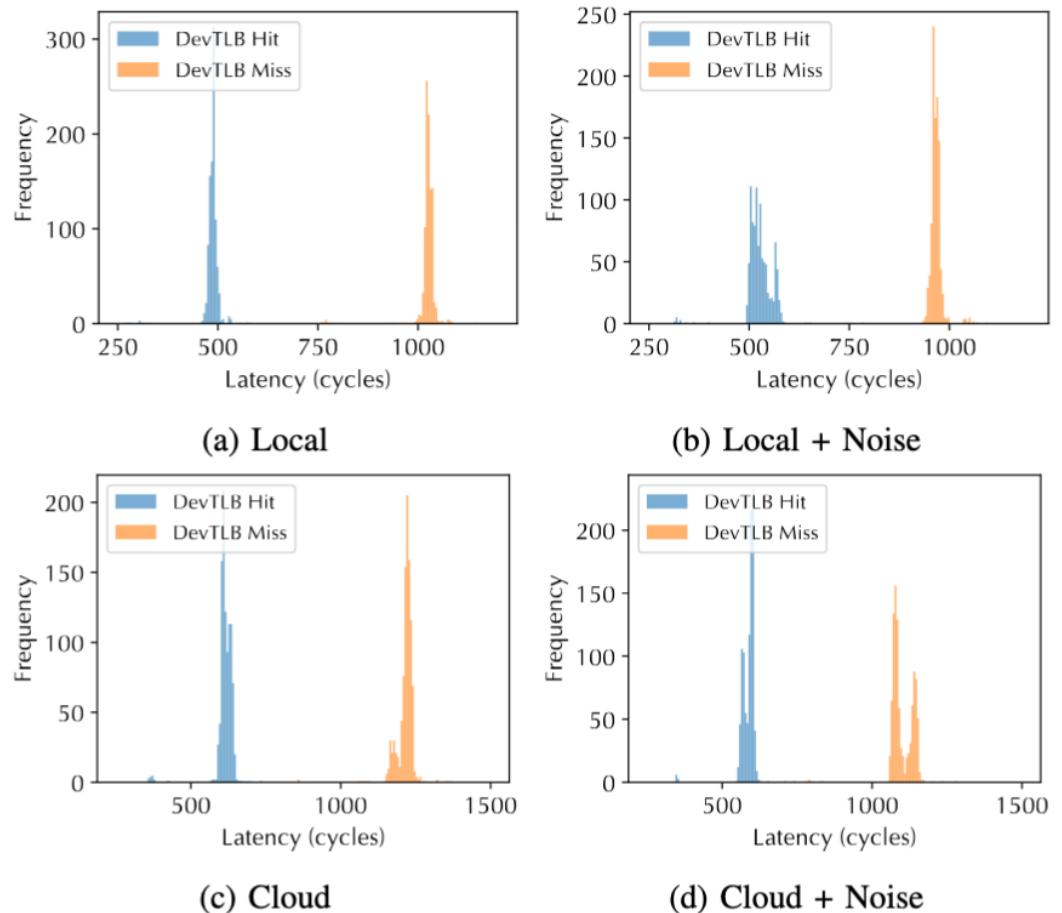
- Redo latency benchmarks on multiple setups
  - Virtualization and PCIe stress
  - Clear impact but still distinguishable
- Attack performance remains
  - Noises outside DSA may not affect the attack effectiveness

	Noisy Local	Cloud	Noisy Cloud	Local + CI
CC <sup>1†</sup>	16.81 / 4.73%	16.97 / 4.69%	16.52 / 5.13%	$17.19 \pm 0.78$
CC <sup>2†</sup>	4.08 / 12.91%	3.96 / 13.3%	3.80 / 13.9%	$4.02 \pm 0.44$
WF	86.0%	85.5%	85.1%	$85.7 \pm 2.8\%$
SSHK <sup>1*</sup>	90.5% / 5.41	93.4% / 5.39	93.0% / 5.35	$5.29 \pm 0.14$
SSHK <sup>2*</sup>	98.2% / 1.25	99.1% / 1.27	98.9% / 1.28	$1.21 \pm 0.09$
LLMC	98.6%	97.7%	98.0%	$98.6 \pm 1.0\%$

<sup>1</sup>  $DSA_{DevTLB}$ . <sup>2</sup>  $DSA_{SWQ}$ .

† format: true capacity (kbps) / bit error rate. CI is in terms of capacity.

\* format:  $F_1$  score / standard deviation (ms). CI is in terms of std.



# Mitigation

## 1 Software Side-Channel Resistance

e.g. Constant-time, data-oblivious execution  
Constantine [CCS'21]: **16% overhead**

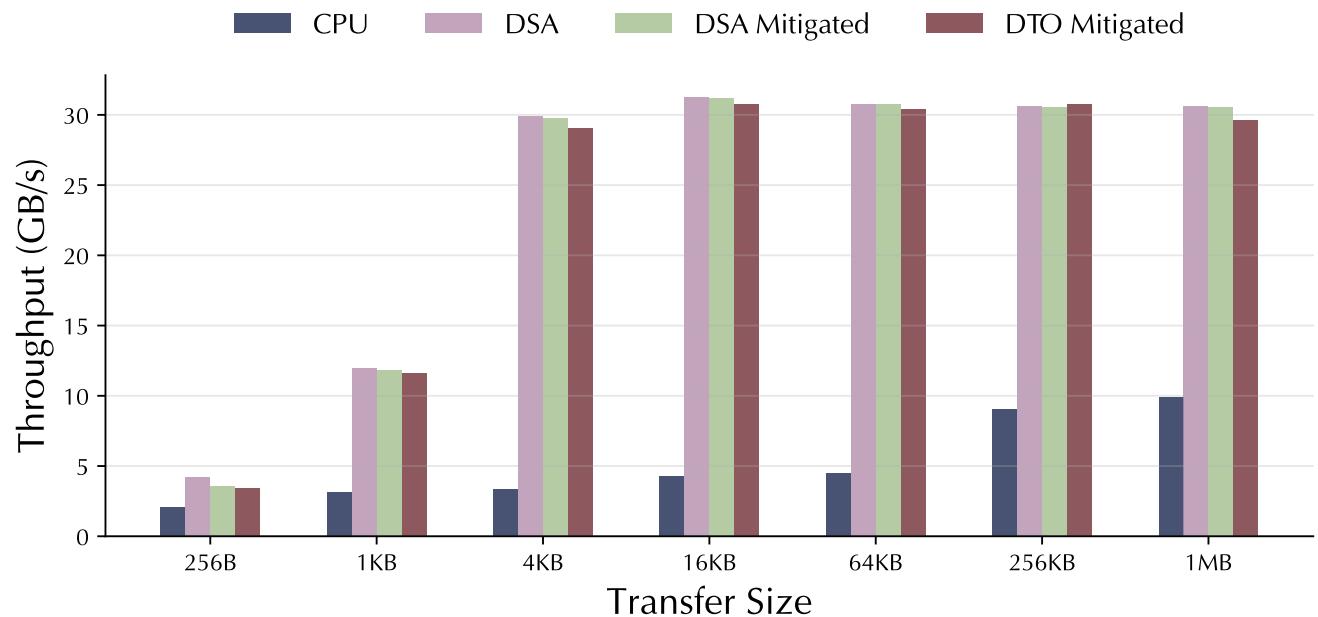
## 3

Our software simulation of DSA's  
**DevTLB partition via flushing**

**Up to 15.7% throughput overhead**  
- Still outperforms CPU's `memcpy`

## 2 Enforcing PASID Isolation

e.g. DevTLB partition  
Secure TLBs [ISCA'19]: **11.4% overhead**



# Conclusion

- We studied Intel DSA and reverse-engineered its microarchitecture
  - DevTLB is indexed by field type and engine ID, and
  - mitigates address translation latency
  - SWQ exposes a timer-less architectural status via DMWr transaction
  - Both are NOT ISOLATED
- We proposed DSASSASSIN, new side-channel attack that
  - bypasses IOMMU
  - can cross VM boundary and is practical in cloud
- We demonstrated the threats by spying on DSA applications
  - Fingerprinting (websites & LLM inference) / Keystroke / Covert Channel

# Thanks for your attention!

## Questions?

DSASSASSIN: Cross-VM Side-Channel Attacks by Exploiting Intel Data Streaming Accelerator  
Ben Chen, Kunlin Li, Shuwen Deng, Dongsheng Wang and Yun Chen. HPCA '26.



香港科技大学(广州)  
THE HONG KONG  
UNIVERSITY OF SCIENCE AND  
TECHNOLOGY (GUANGZHOU)



清华大学