

Perspective: A Principled Framework for Pliable and Secure Speculation in Operating Systems

Tae Hoon Kim, David Rudo, Kaiyang Zhao, Zirui Neil Zhao[†], Dimitrios Skarlatos

Carnegie Mellon University

[†]University of Illinois Urbana-Champaign

[†]The University of Texas at Austin (New Affil.)

ISCA '24 – Session 6B: Security

Spectre-V1: A Quick Recap



Adversary gives malicious
input of $x > \text{size}$

```
if (x < size) { // Mispredicted branch
    int secret = array1[x]; // access
    int y = array2[secret * 4096]; // transmit
}
```

Speculative out-of-bound access

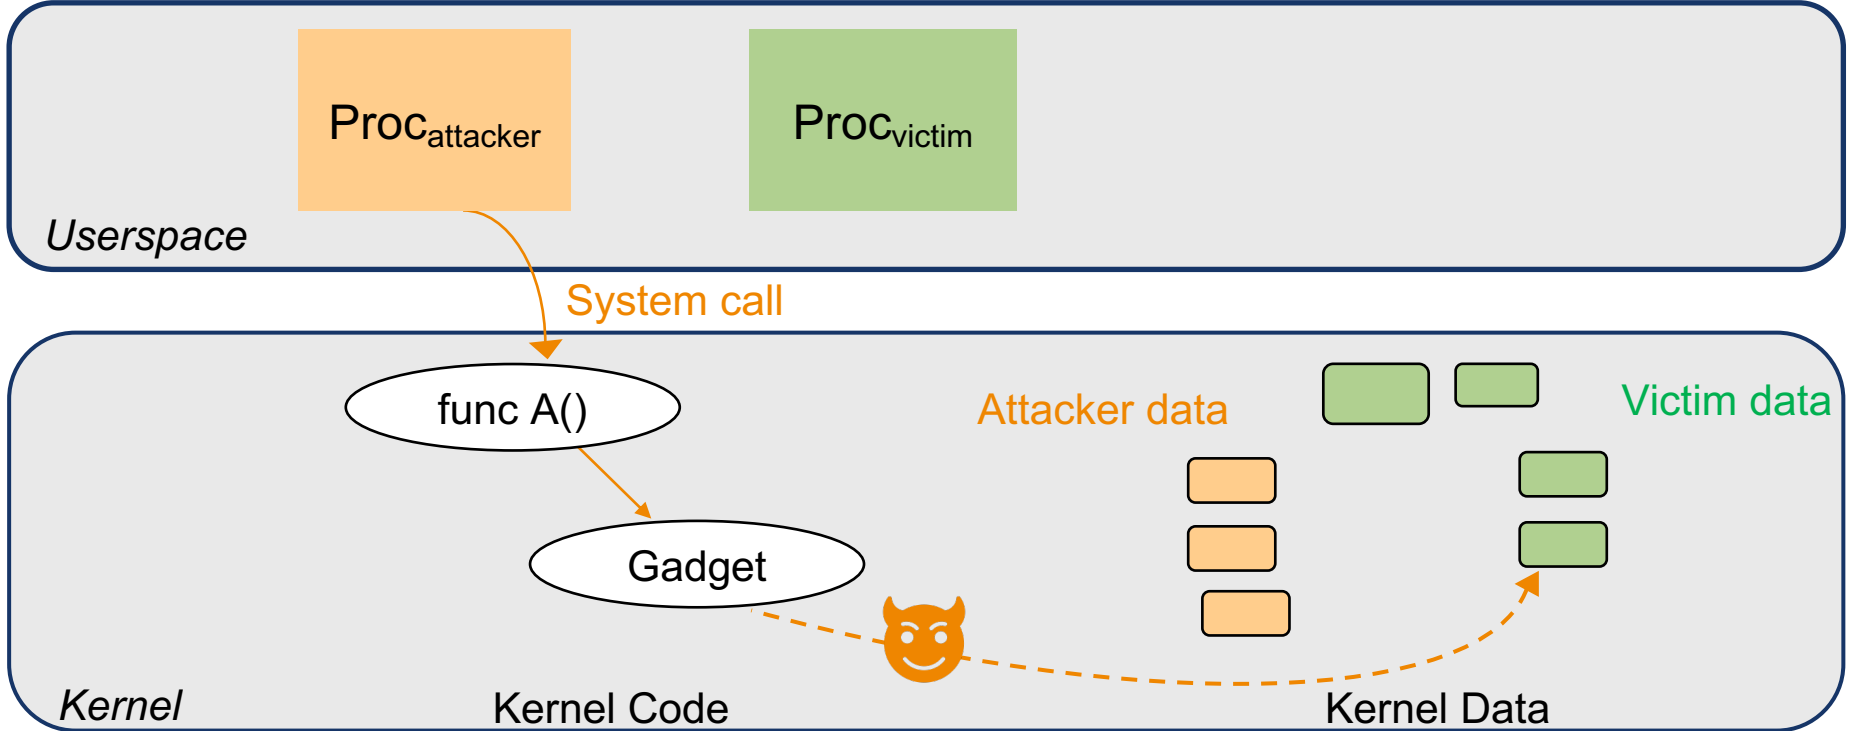


Spectre Gadget



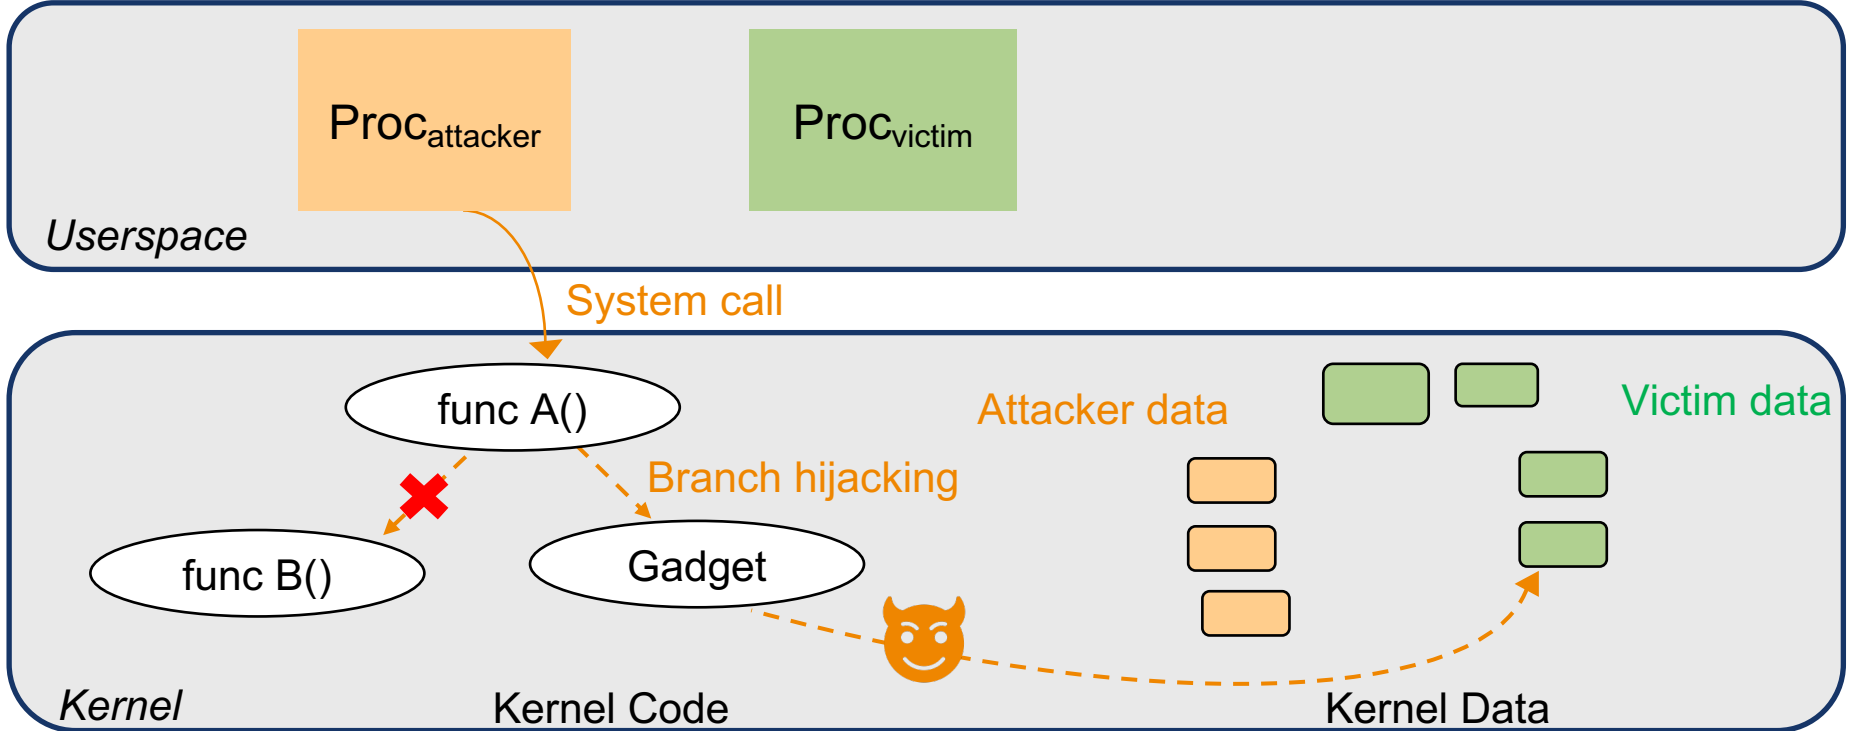
Spectre-V1: Gadget in the Kernel

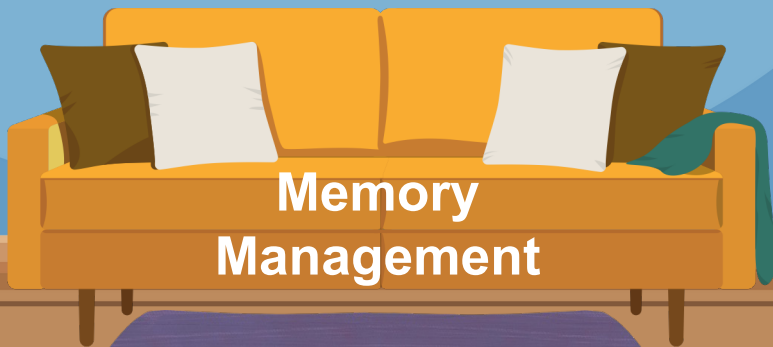
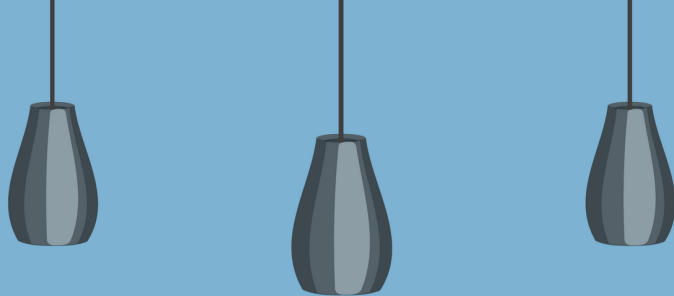
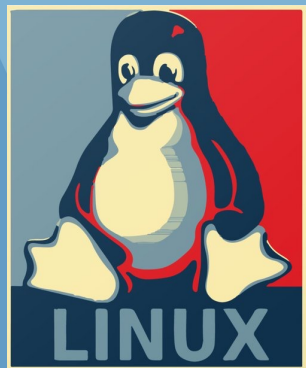
Transient 
Non-transient 



Spectre-V2: Gadget in the Kernel

Transient 
Non-transient 

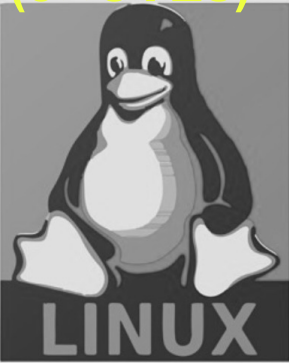




Networking



Branch
History
Injection
(3×CVEs)



Retbleed (2×CVEs)



Storage



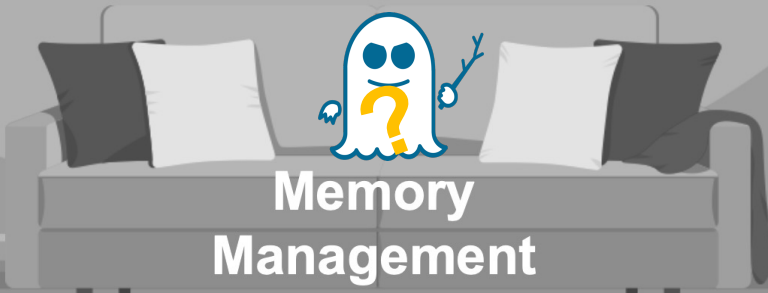
Misuse of mitigations (≥4×CVEs)

LFENCE/JMP on AMD (1×CVE)



Missing
Mitigation
(1×CVE)

KVM



Memory
Management



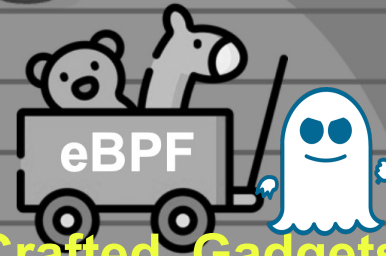
ptrace
Gadget by Accident
(1×CVE)



Networking

Gadget in Xilinx Driver
(1×CVE)

Device Drivers



User-Crafted Gadgets

(≥ 6×CVEs)

Existing Defenses are Software Agnostic

```
if (x < size) { // Unresolved branch
```

```
    int secret = array1[x]; // access
```

```
    int y = array2[secret * 4096]; // transmit
```

```
}
```

Block or protect speculative accesses (SLH, Fence, InvisiSpec, DOM, ...)

Block speculative transmissions (STT, NDA, ...)



Existing Defenses → Complex HW + Over-Protection

Existing Defenses are software agnostic
But these solutions introduce complex hardware ☹️
Over-protect with high performance overhead

```
if (x < size) { // Unresolved branch  
    int secret = array1[x]; // access  
    int y = array2[secret * 4096]; // transmit  
}
```

Block or protect speculative accesses (SLH, Fence, InvisiSpec, DOM, ...)

Block speculative transmissions (STT, NDA, ...)



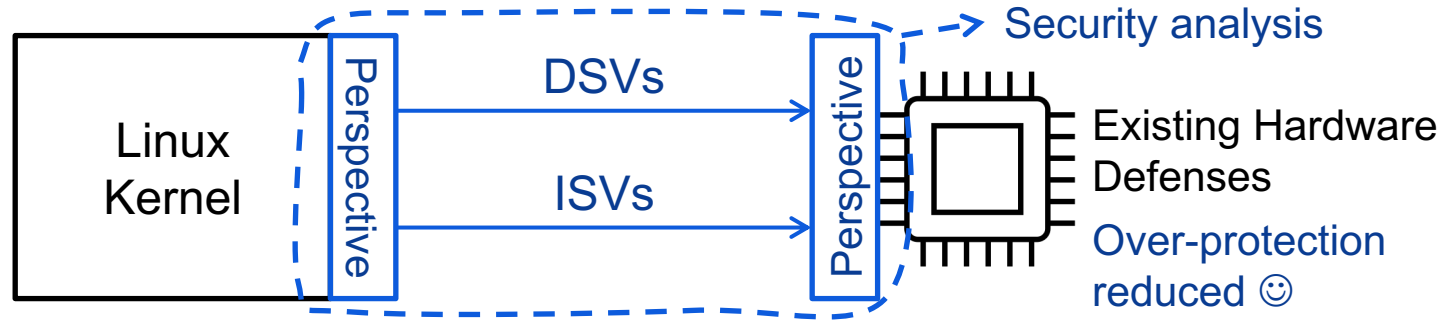
Perspective: An OS-HW Interface for Secure Speculation in OS

- **A taxonomy of transient execution attacks in the kernel**

Active and passive attack scenarios ⇒ Agnostic to attack variants (e.g., Spectre v1, v2, ...)

- **Perspective OS-HW interface ⇒ Tailored defense**

- Data speculation views (DSVs): which *data* can be speculatively accessed
- Instruction speculation views (ISVs): which *instructions* can speculatively execute
- Both views can be associated with an execution context (e.g., a process)



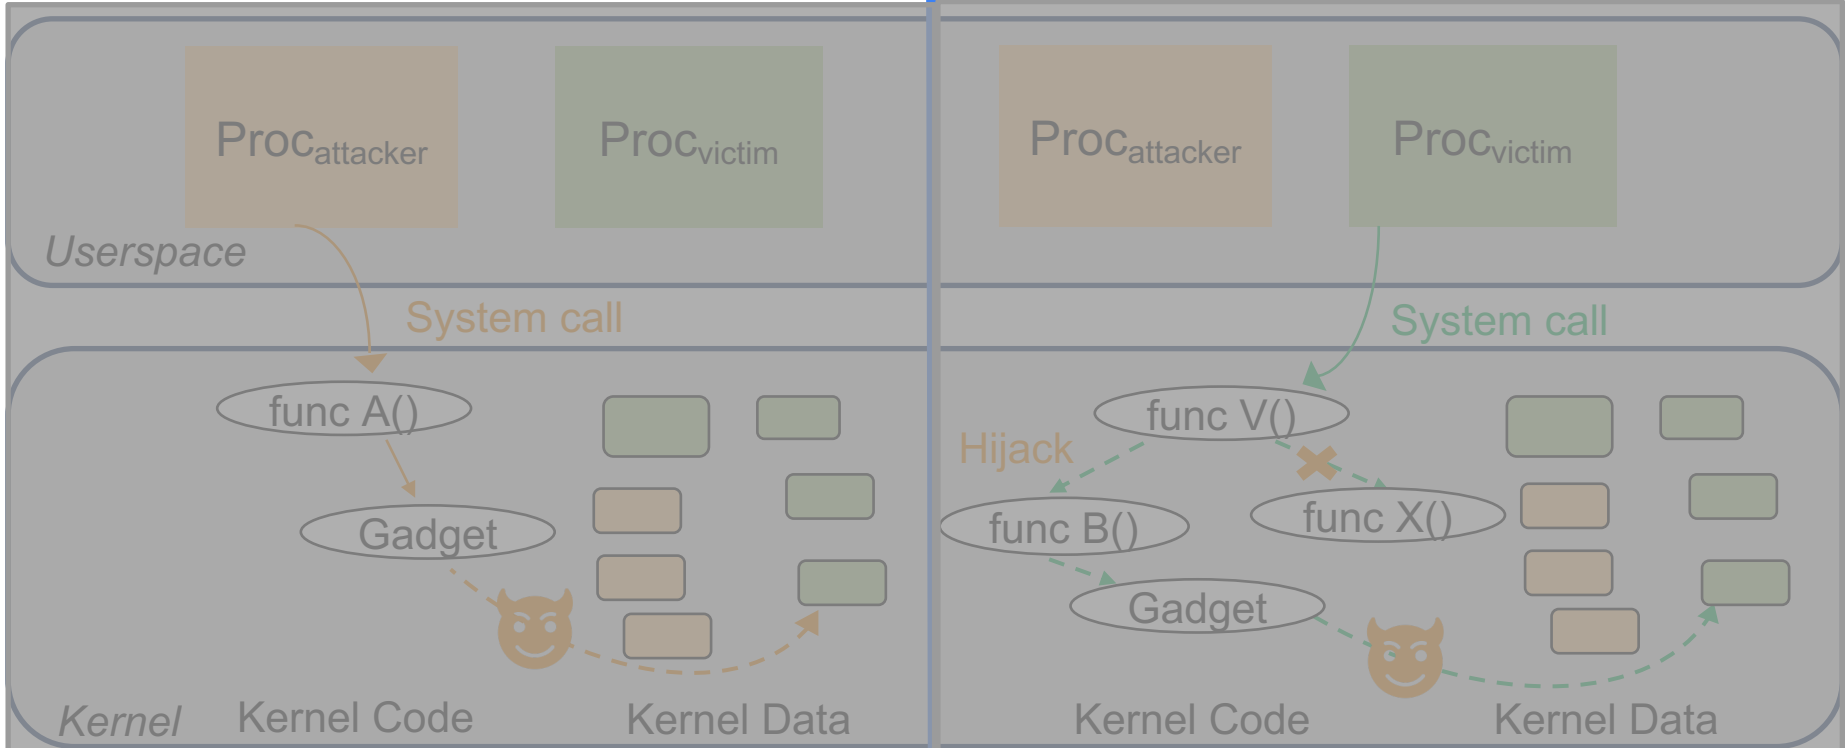
Perspective's **tailored protection** allows the use of **simple** hardware defenses like Fence while maintaining a **low execution overhead** (e.g., 1.2% on datacenter applications)

Kernel Gadget Taxonomy

Transient 
Non-transient 

Active Attacks

Passive Attacks

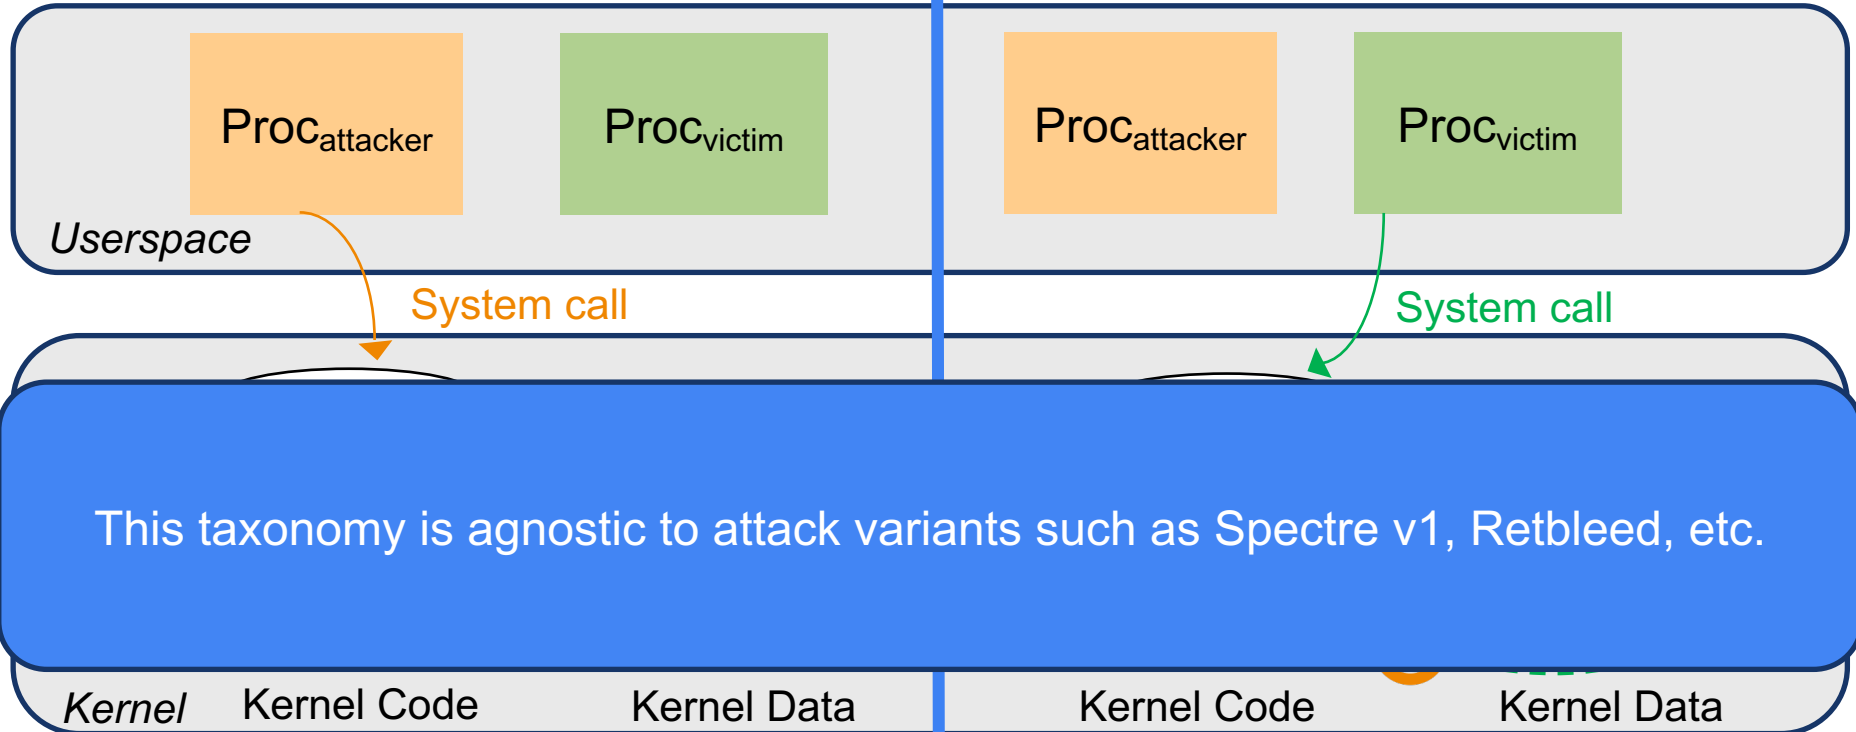


Kernel Gadget Taxonomy

Transient ----->
Non-transient ----->

Active Attacks

Passive Attacks

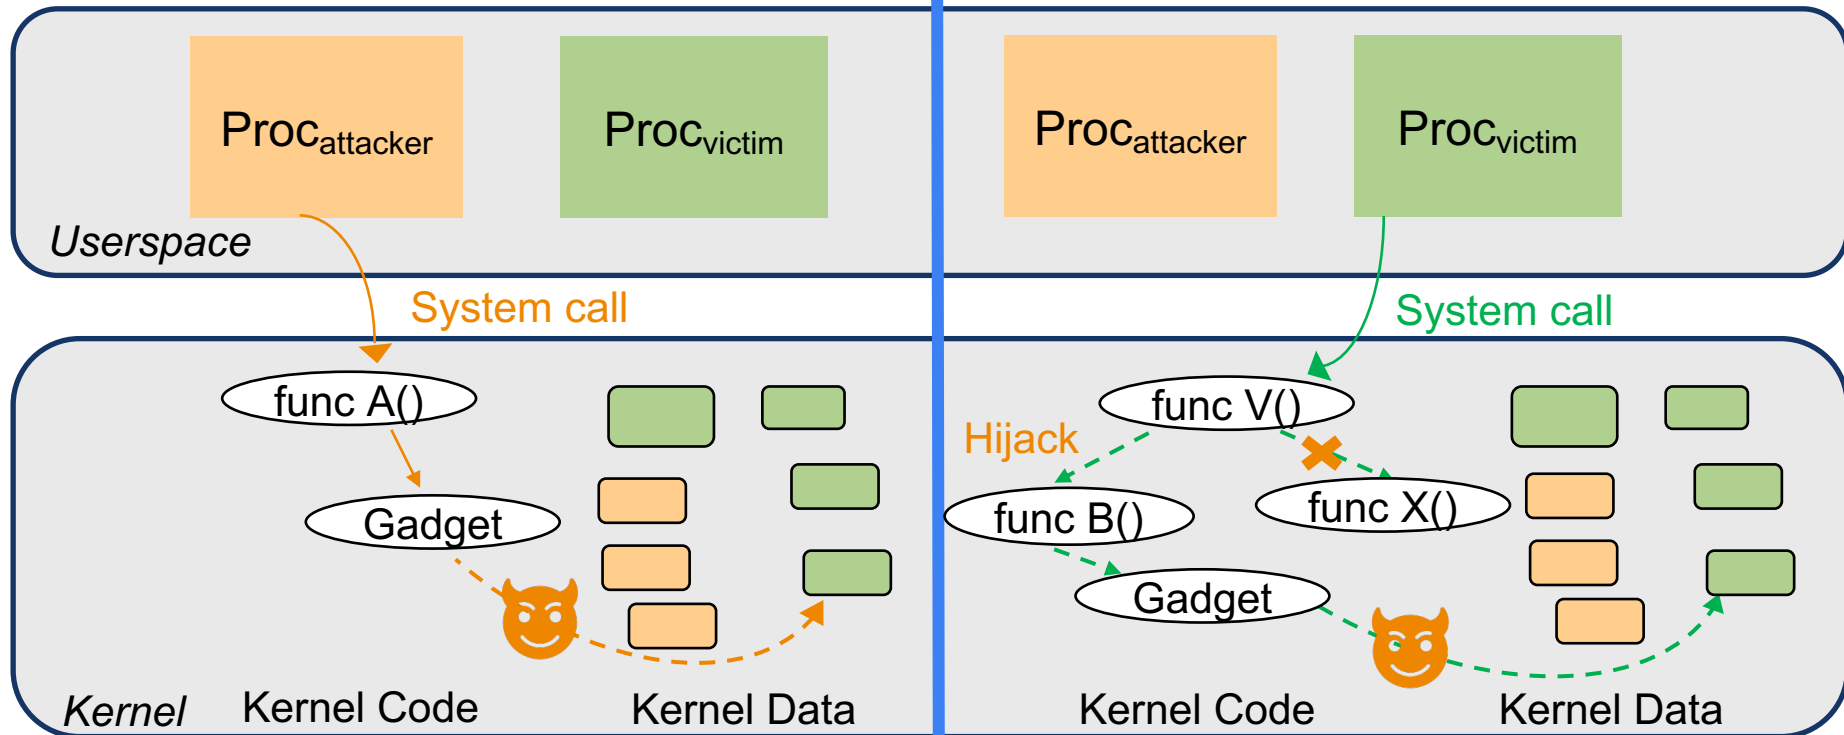


Perspective: Data Speculation Views

Transient 
Non-transient 

Active Attacks

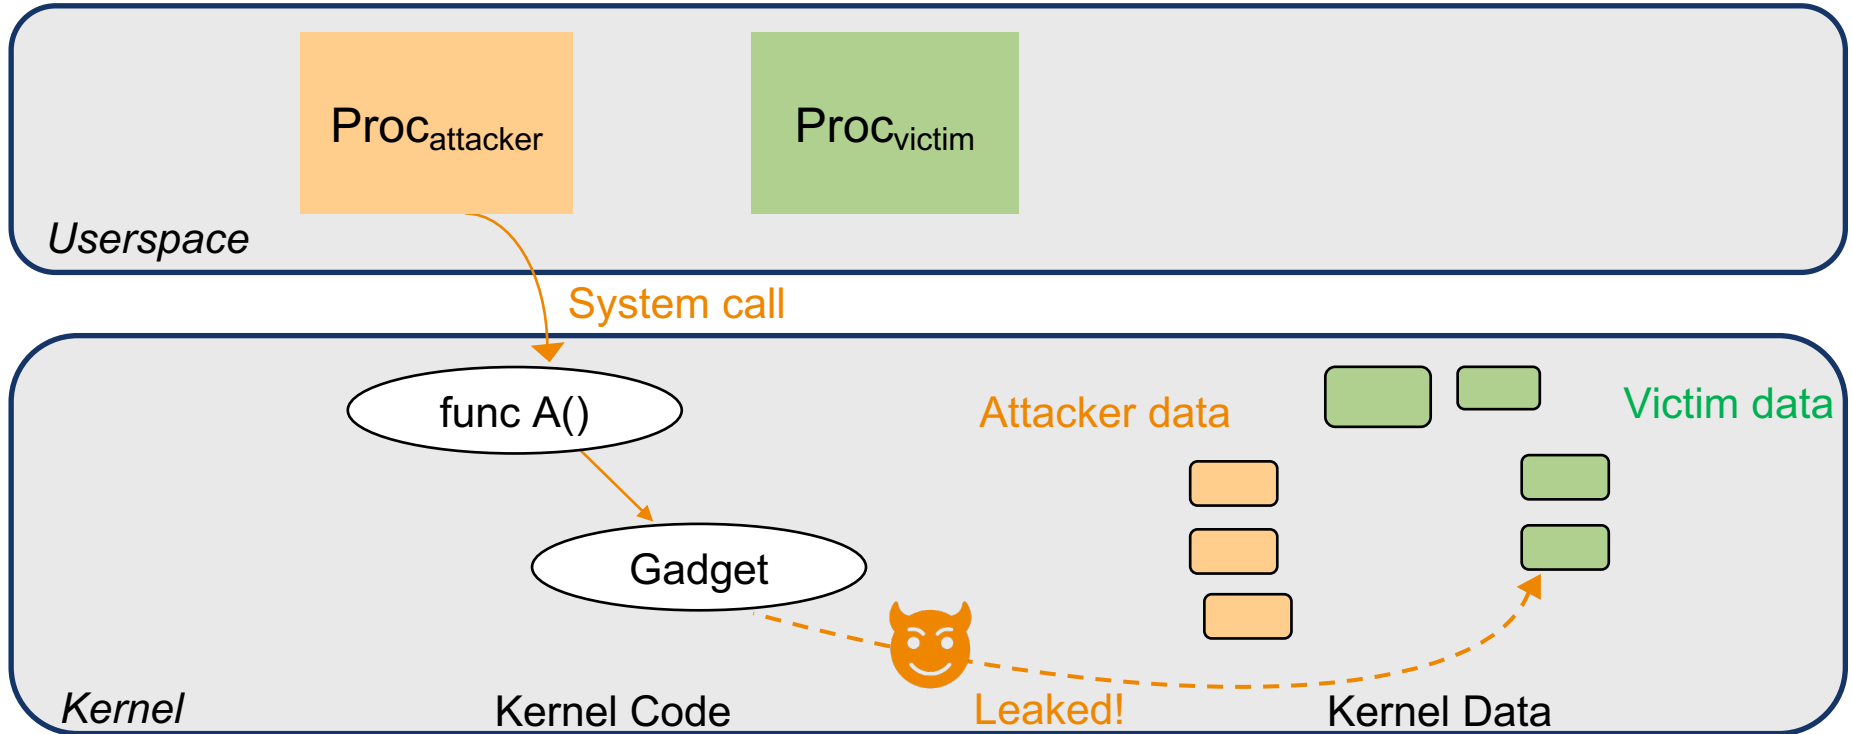
Passive Attacks



Perspective: Data Speculation Views

Active Attacks

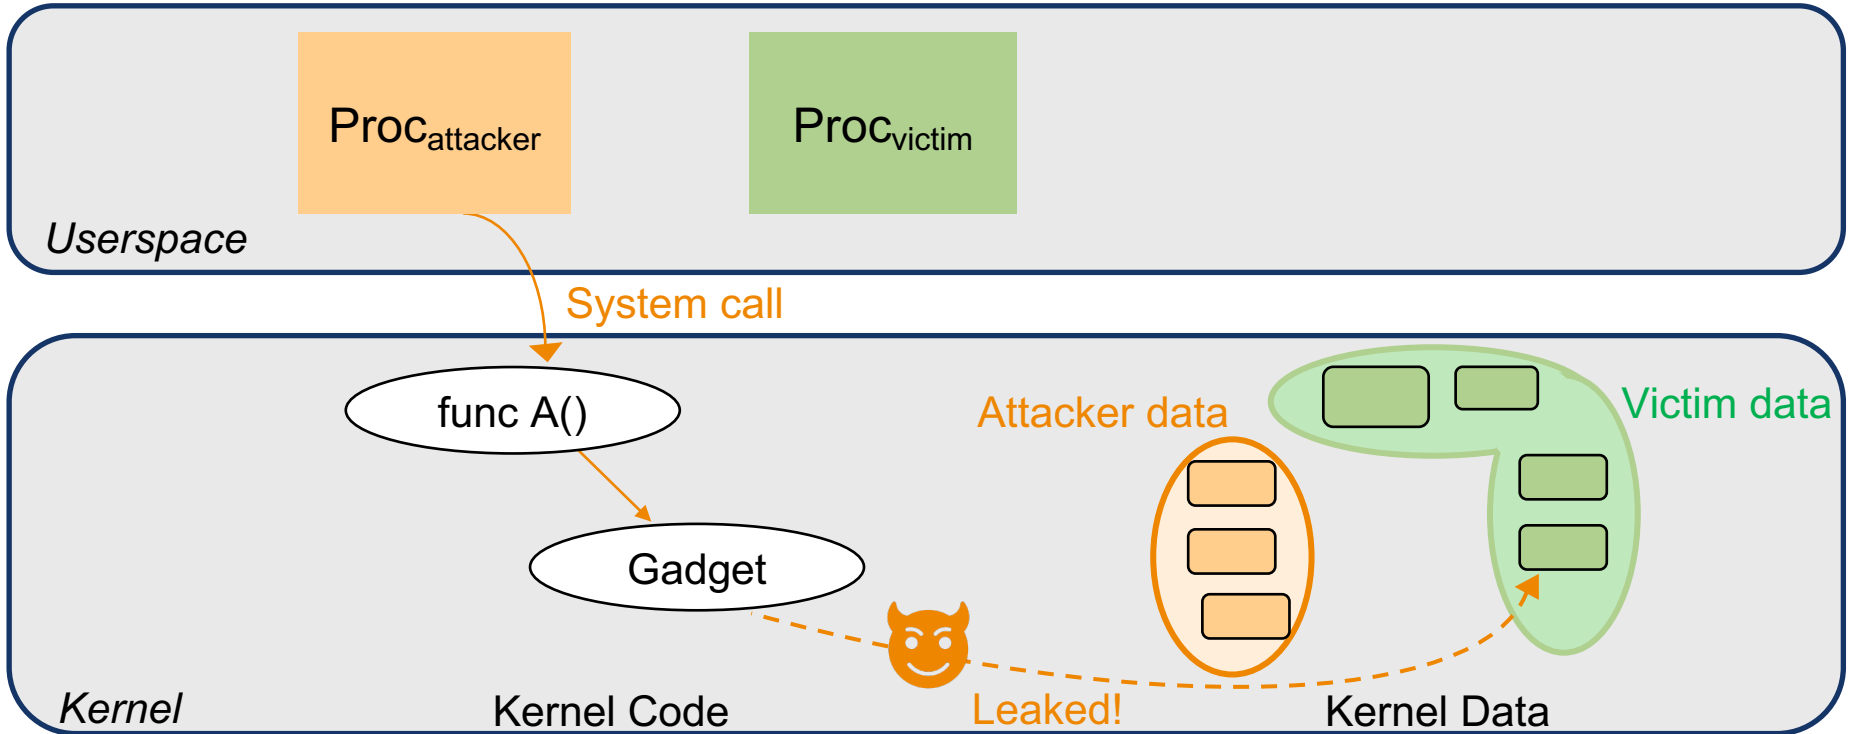
Transient 
Non-transient 



Perspective: Data Speculation Views

Active Attacks

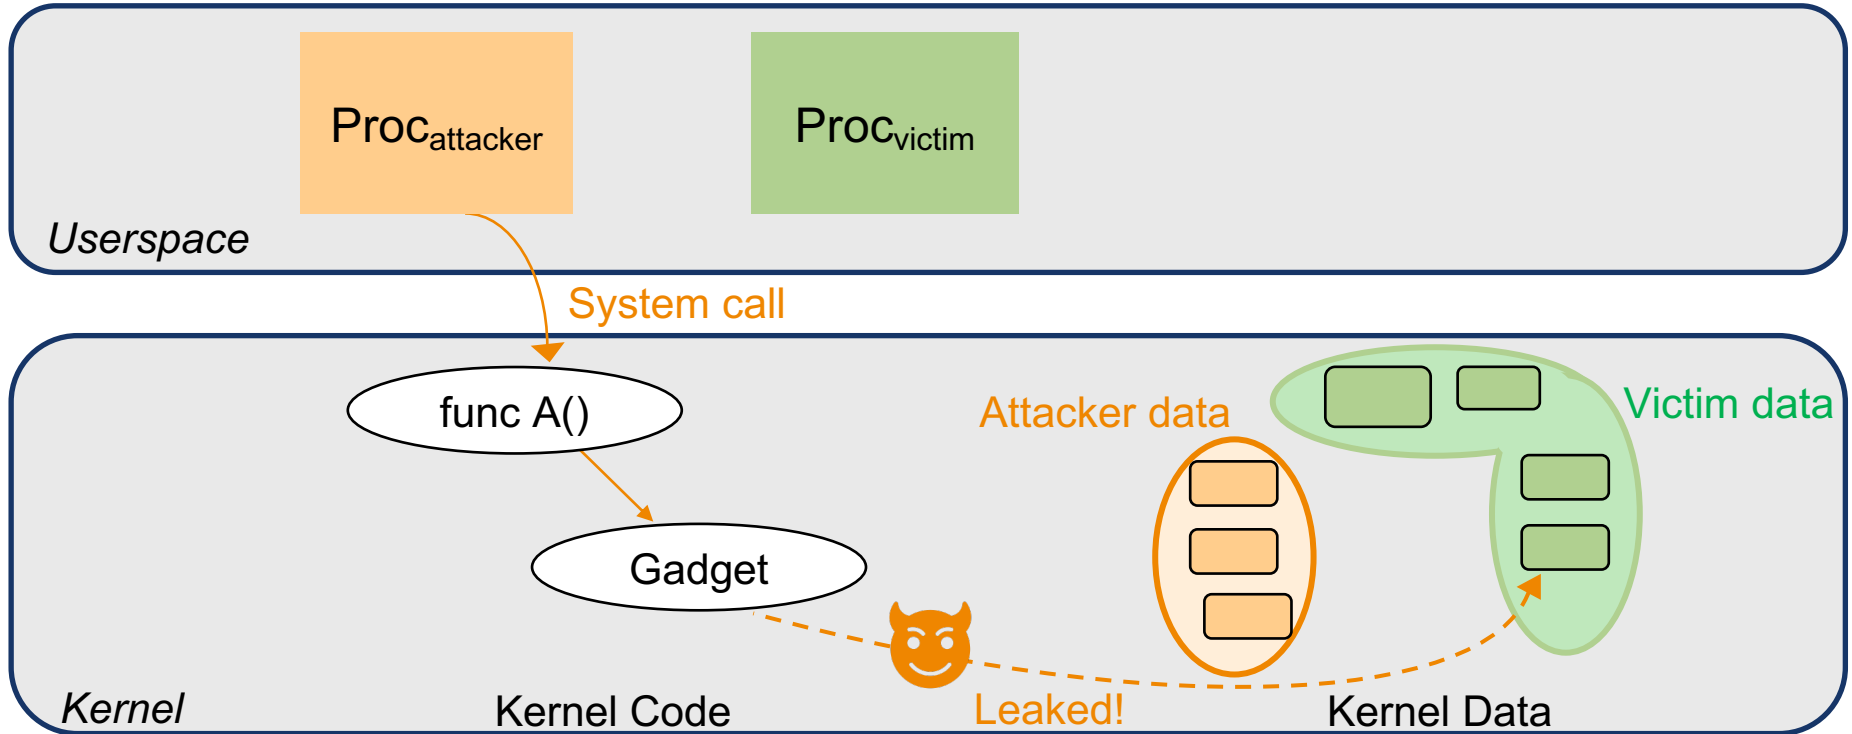
Transient 
Non-transient 



Perspective: Data Speculation Views

Active Attacks

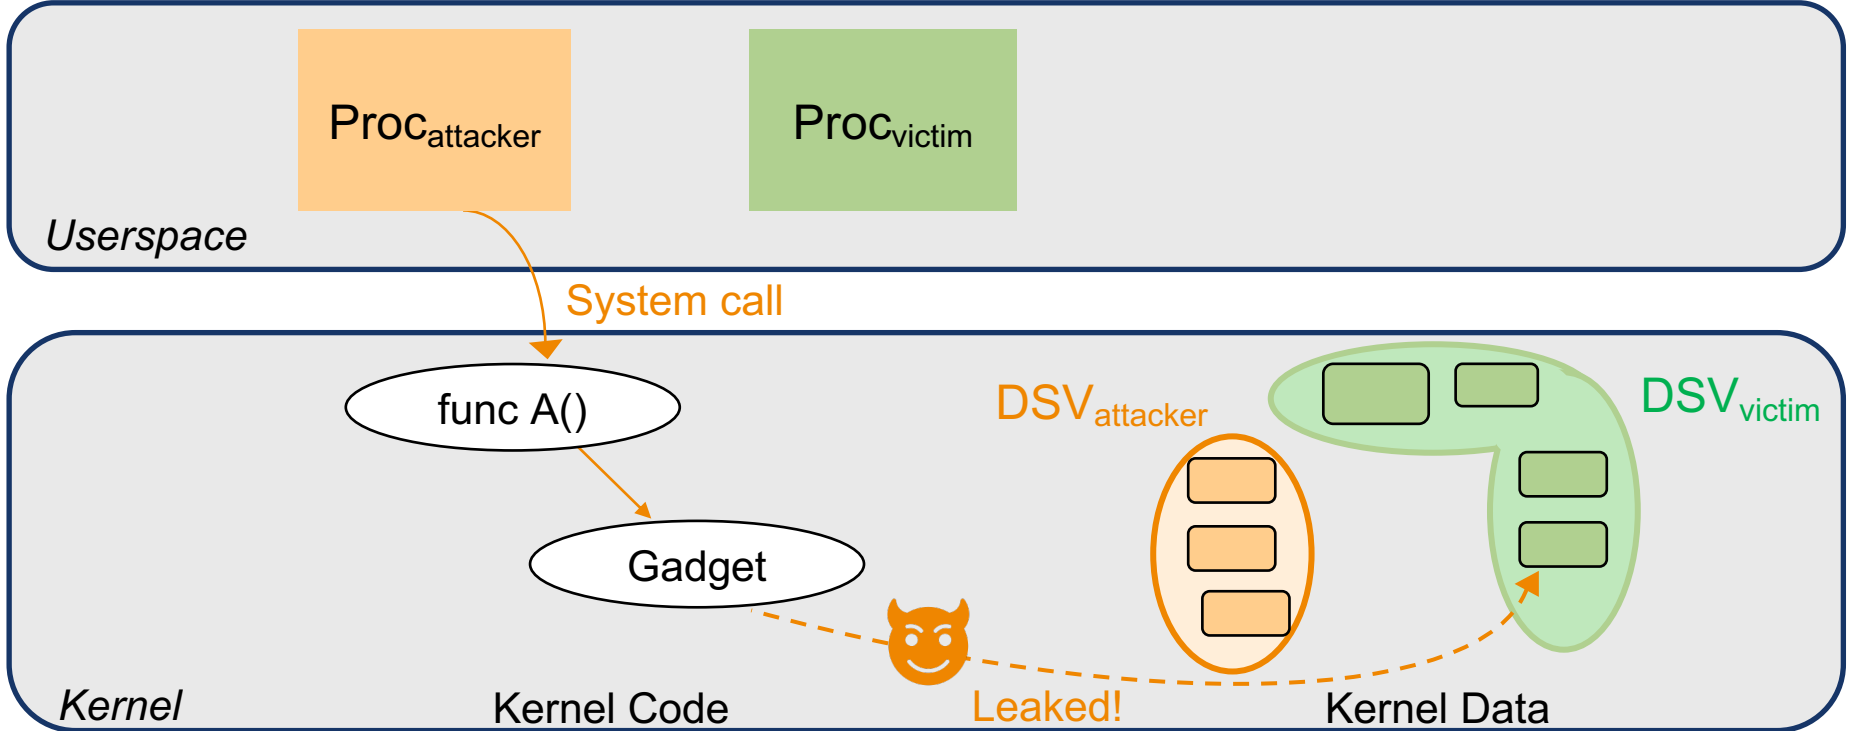
Transient 
Non-transient 



Perspective: Data Speculation Views

Active Attacks

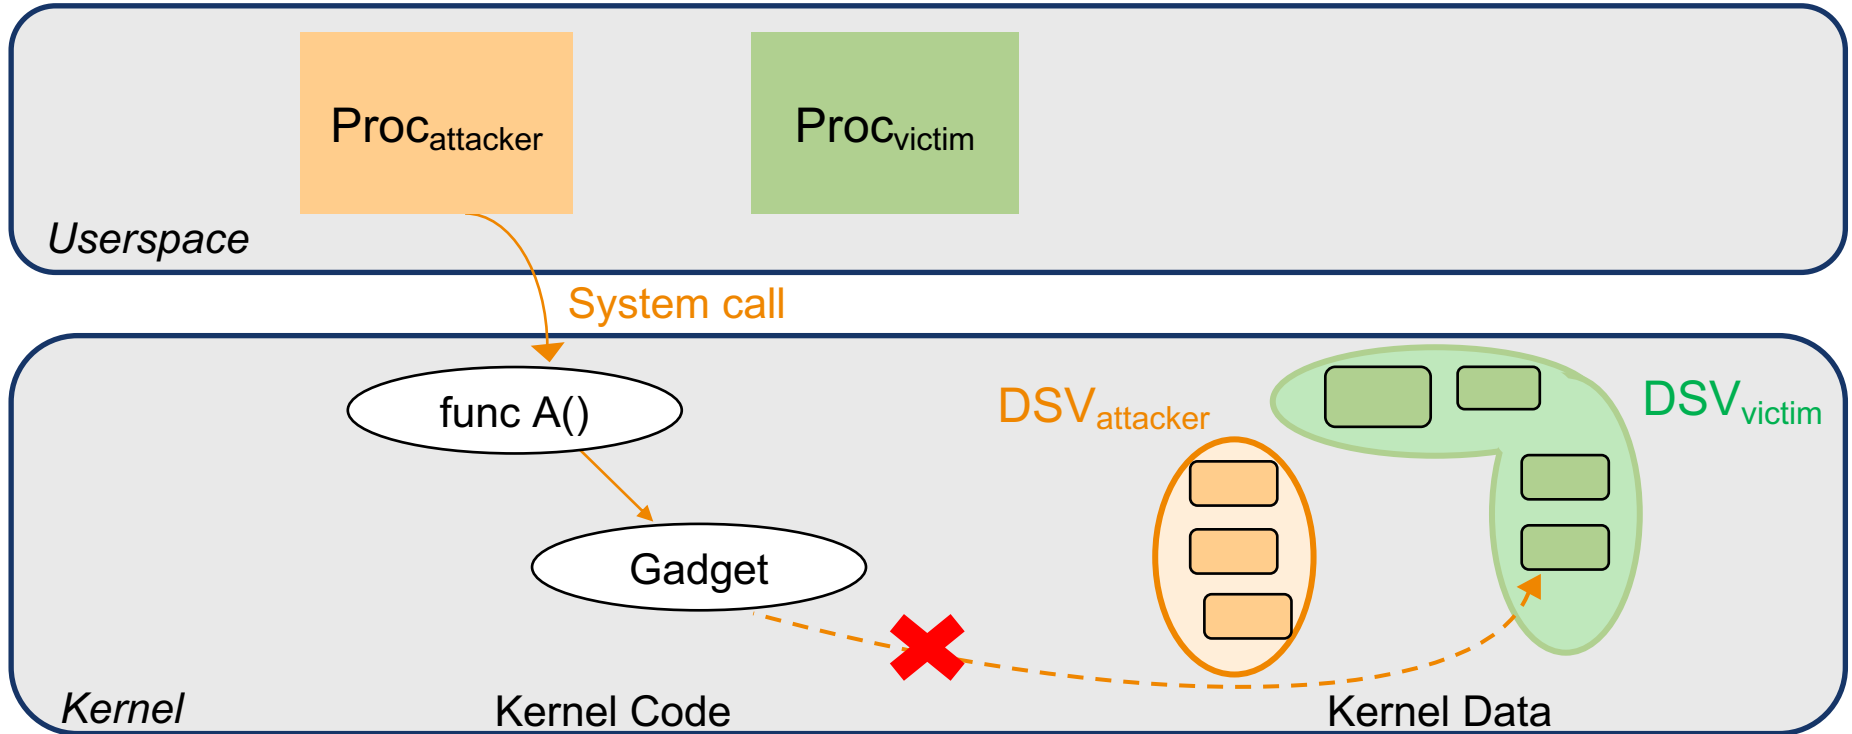
Transient 
Non-transient 



Perspective: Data Speculation Views

Active Attacks

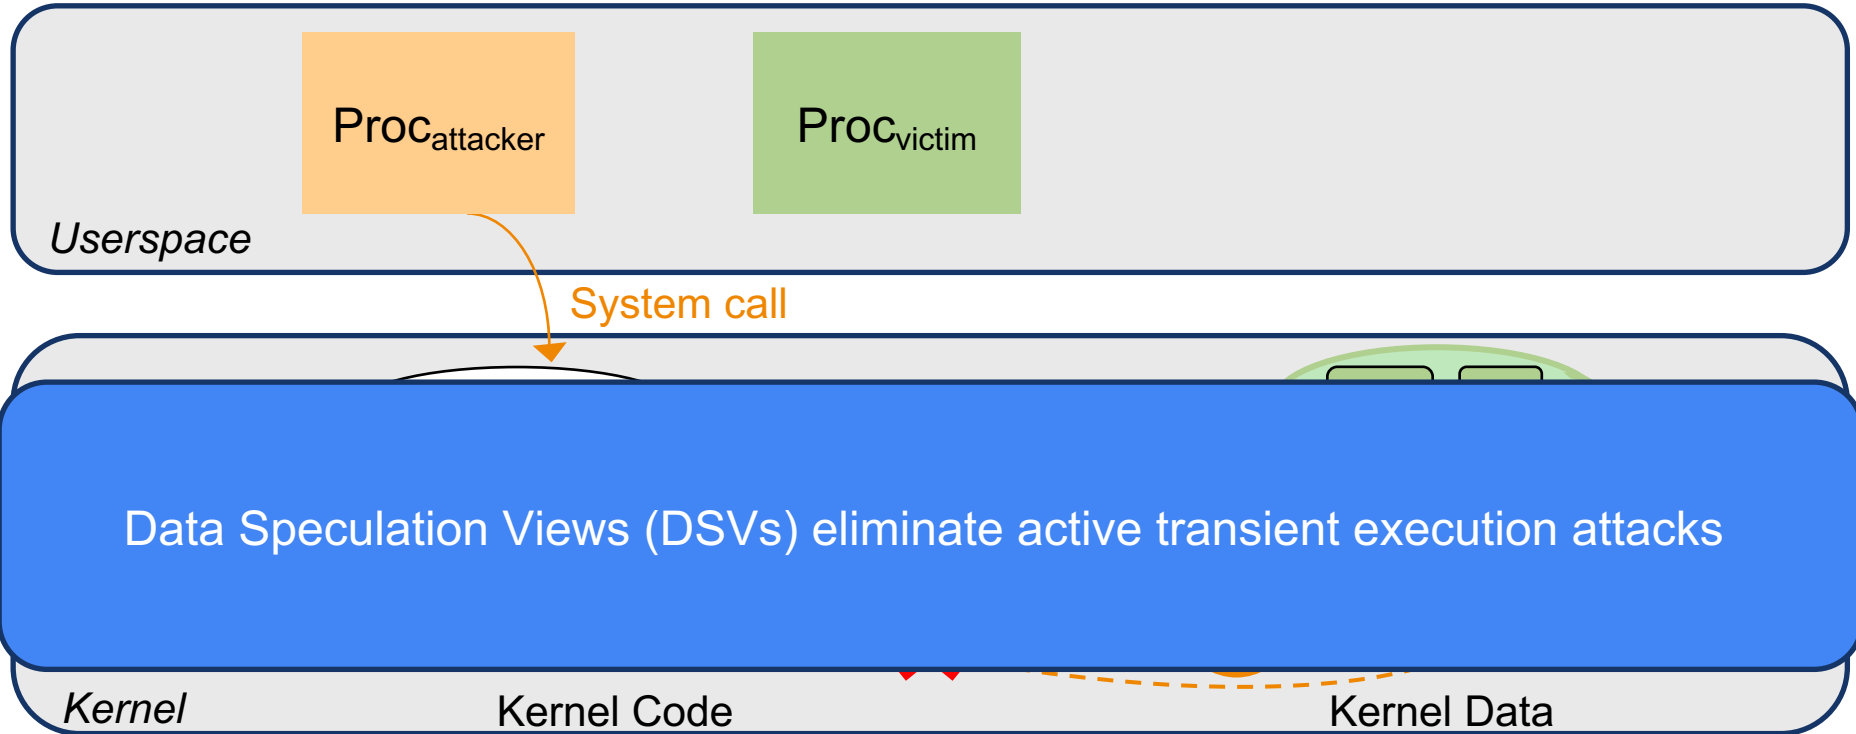
Transient \dashrightarrow
Non-transient \longrightarrow



Perspective: Data Speculation Views

Transient - - - ->
Non-transient — — —>

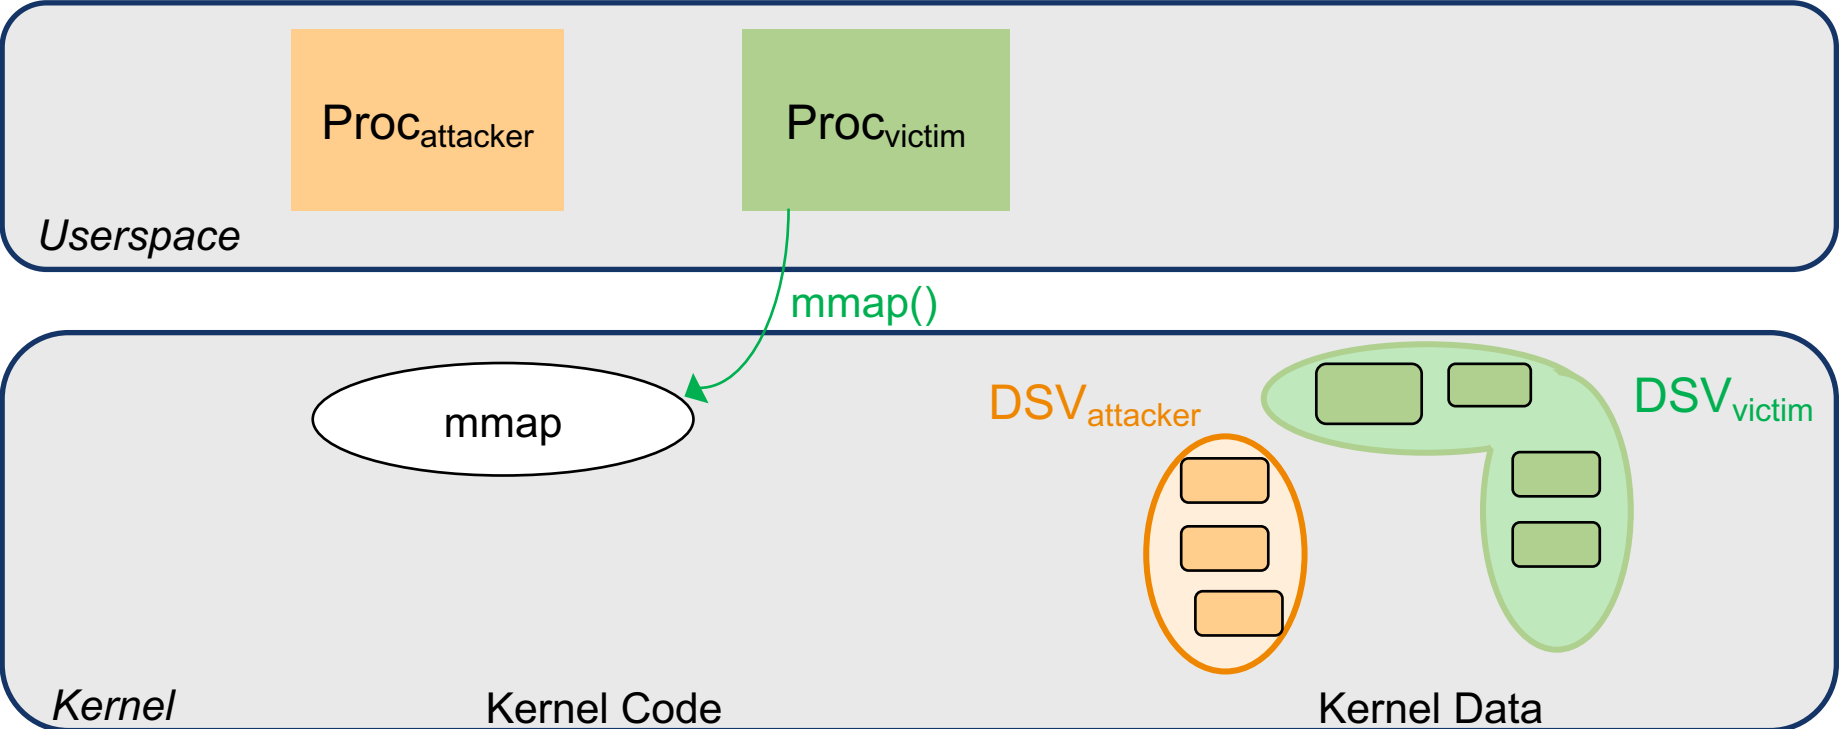
Active Attacks



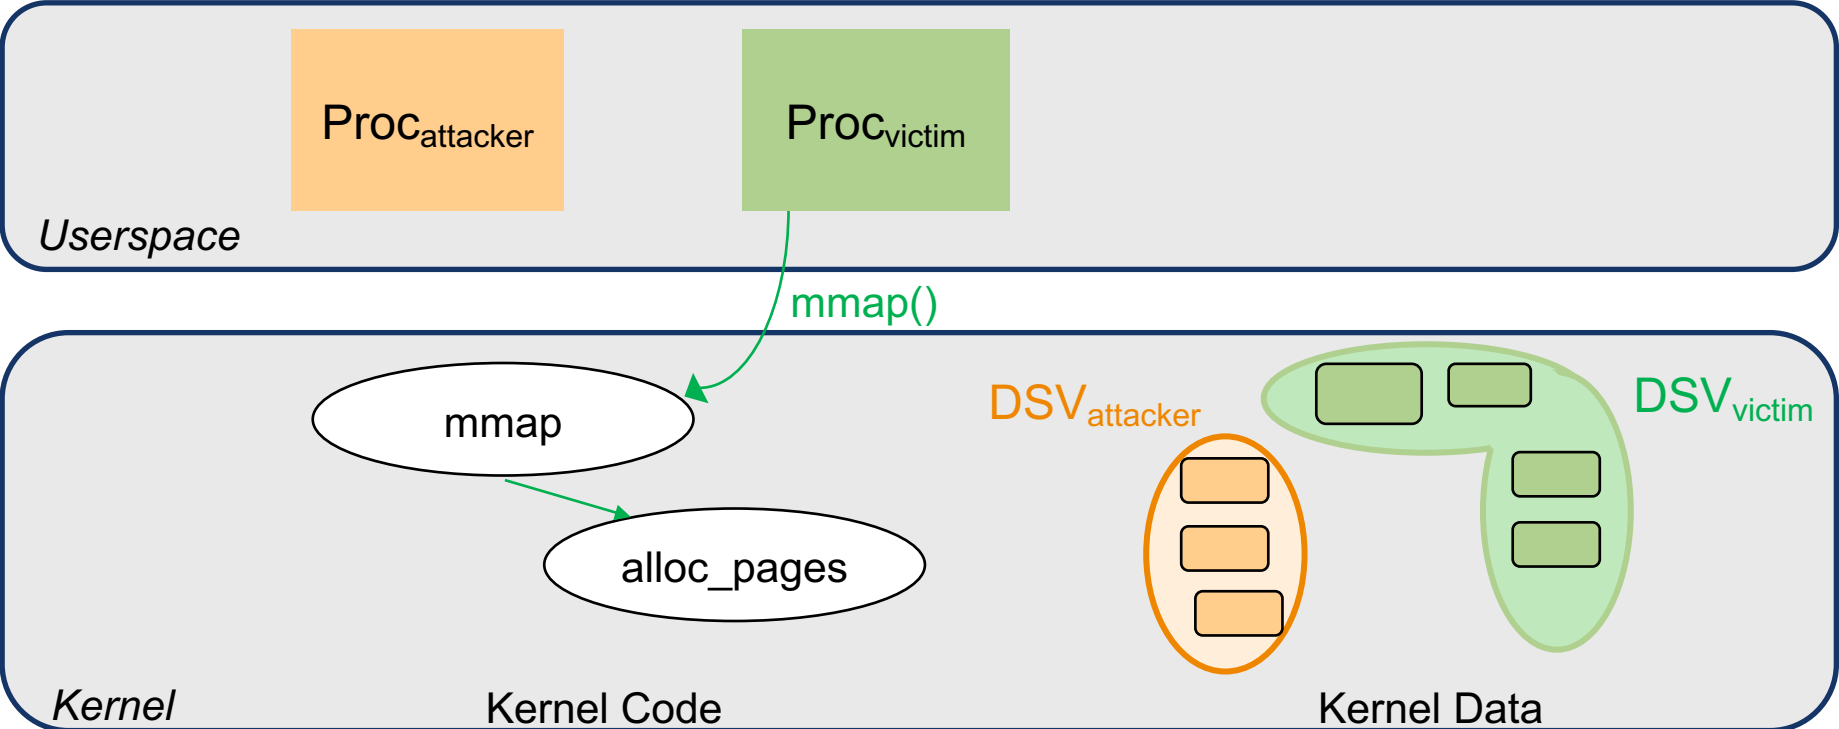
DSV Through Allocation



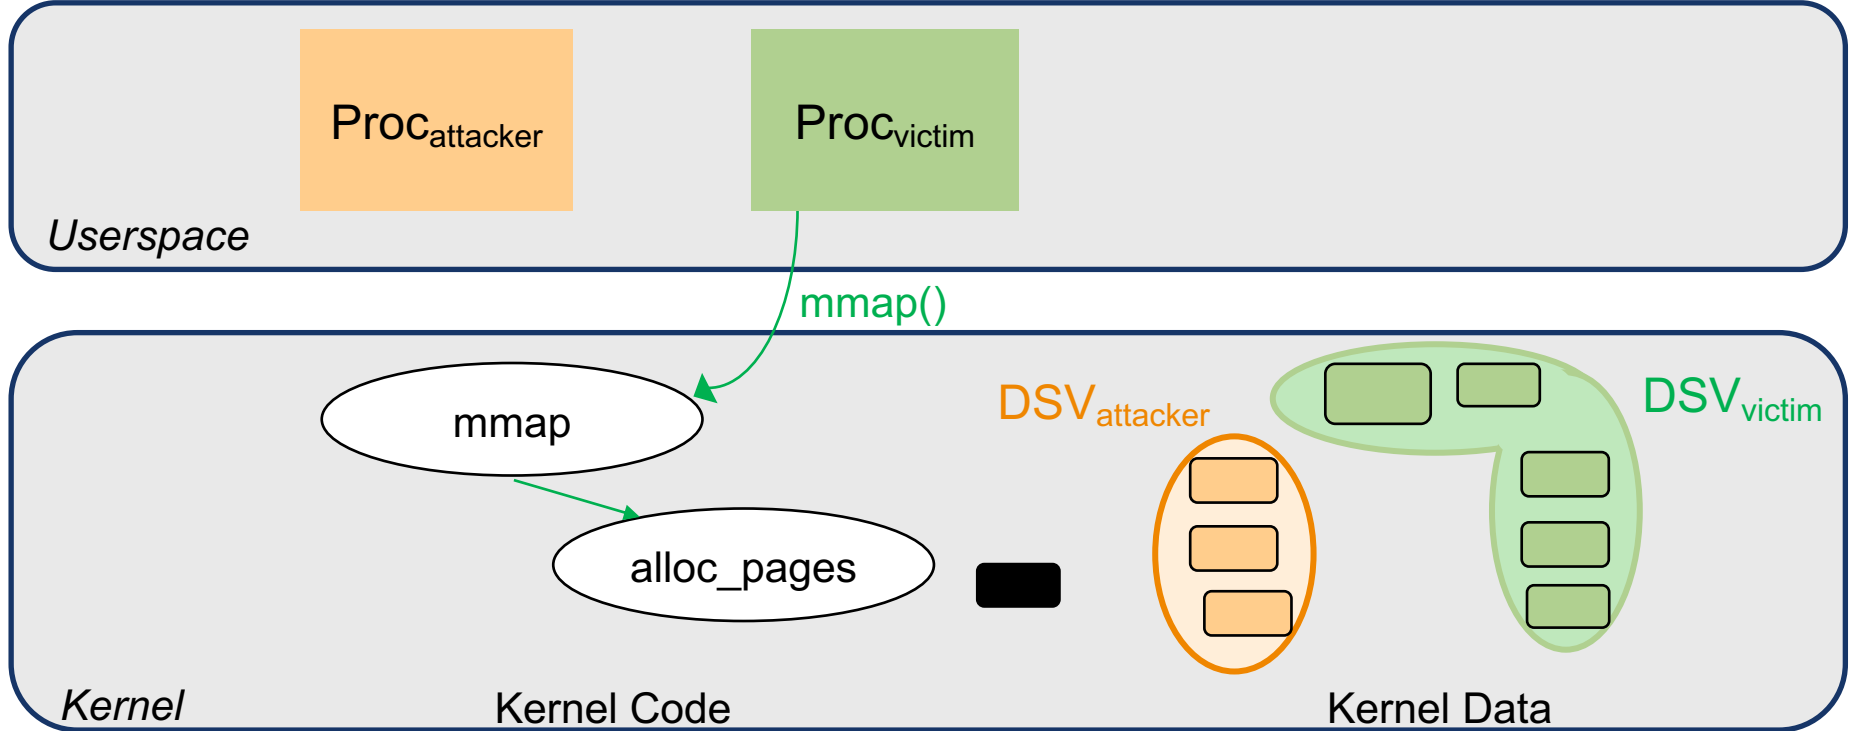
DSV Through Allocation



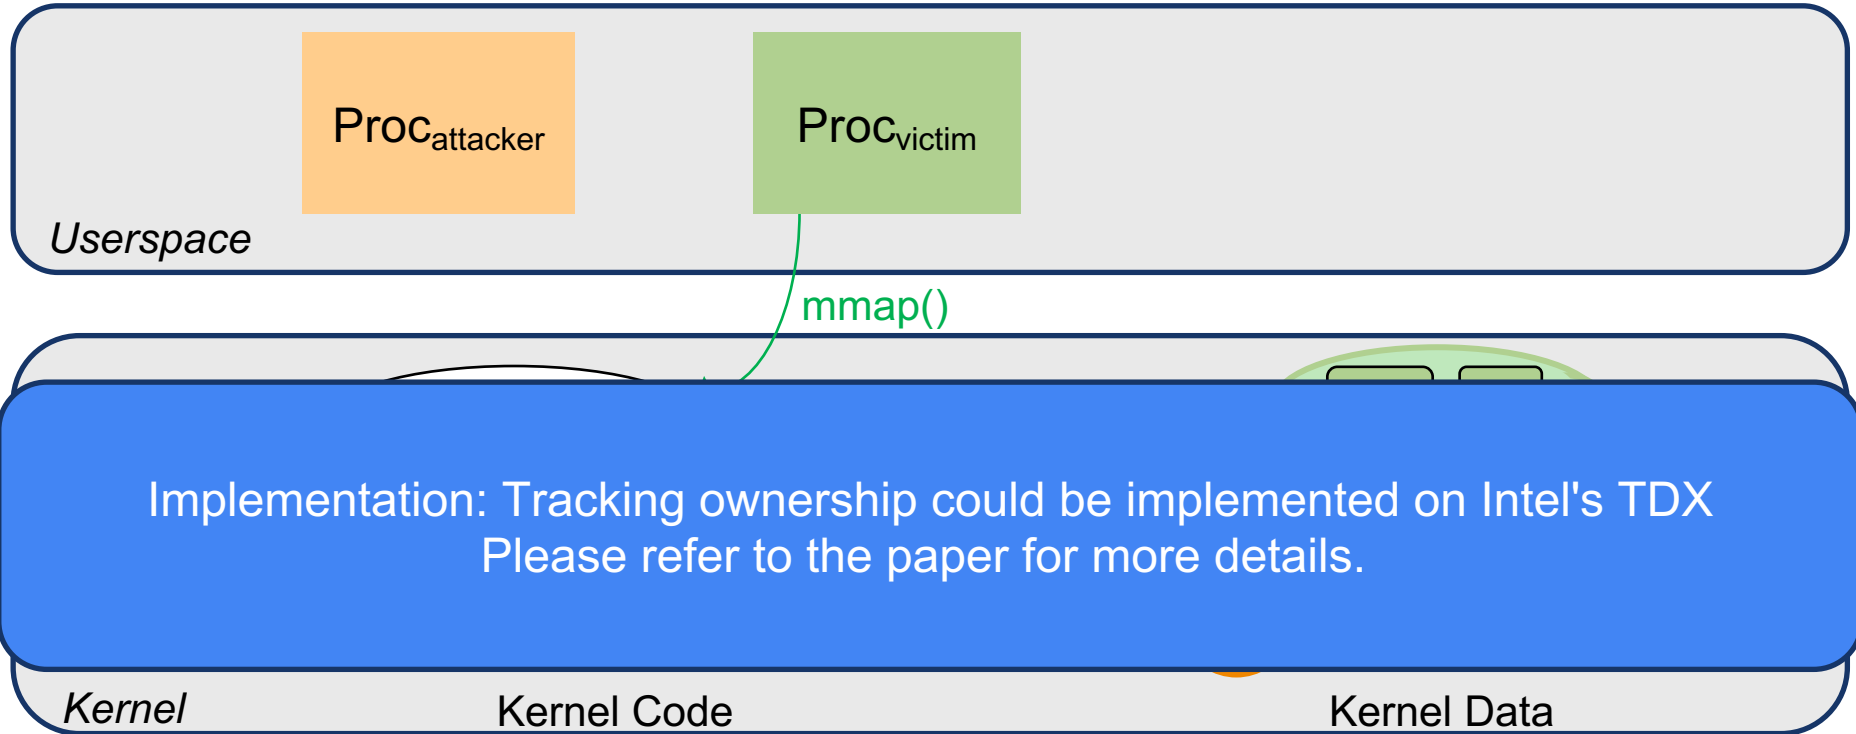
DSV Through Allocation



DSV Through Allocation



DSV Through Allocation

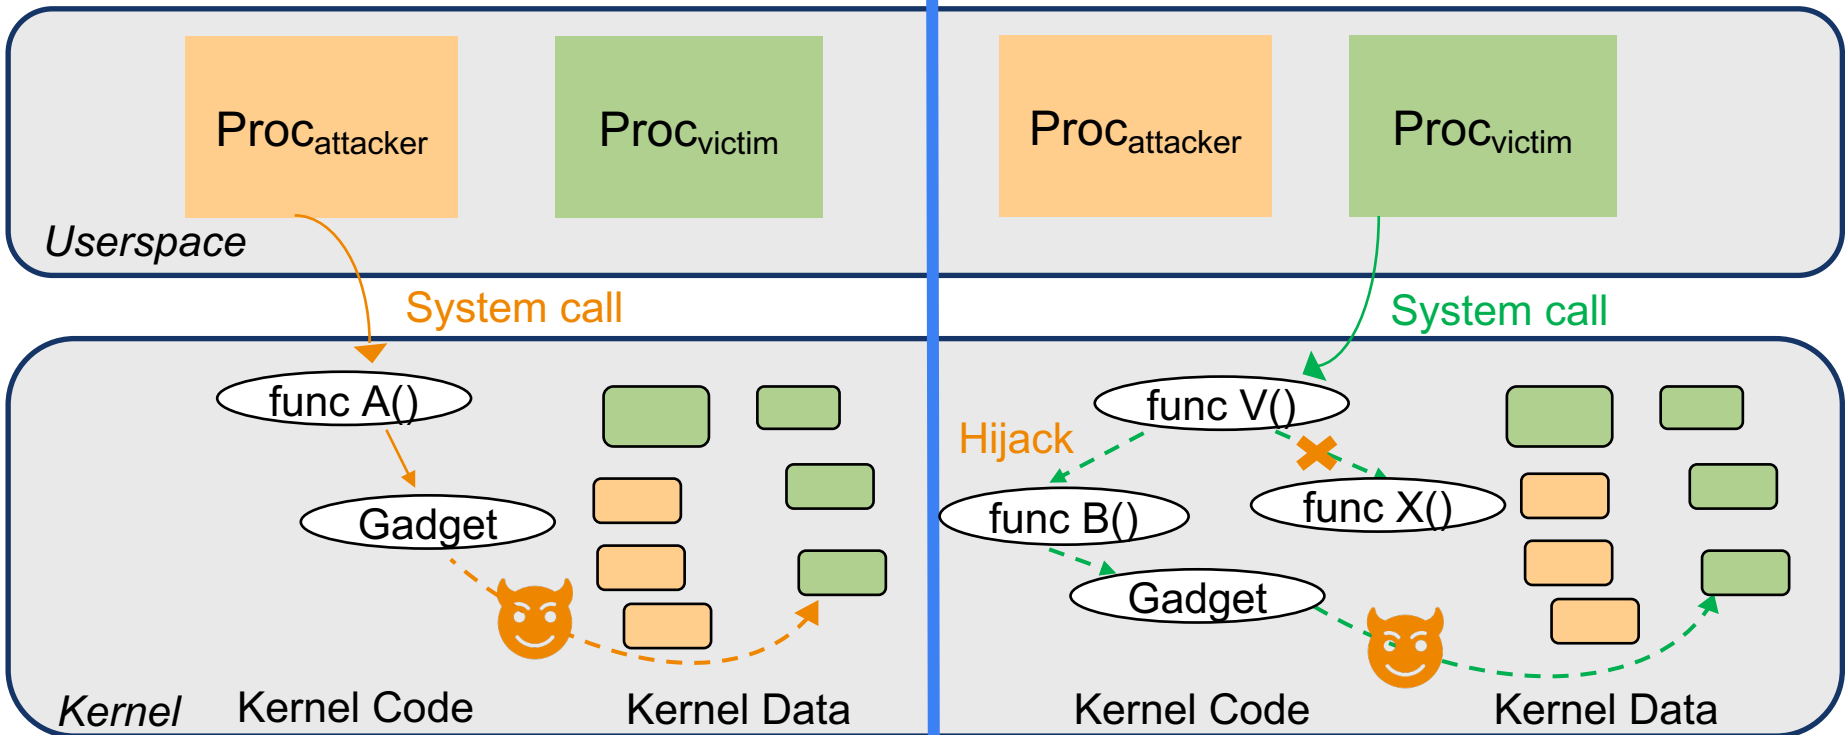


Perspective: Instruction Speculation Views

Transient ---→
 Non-transient —→

Active Attacks

Passive Attacks

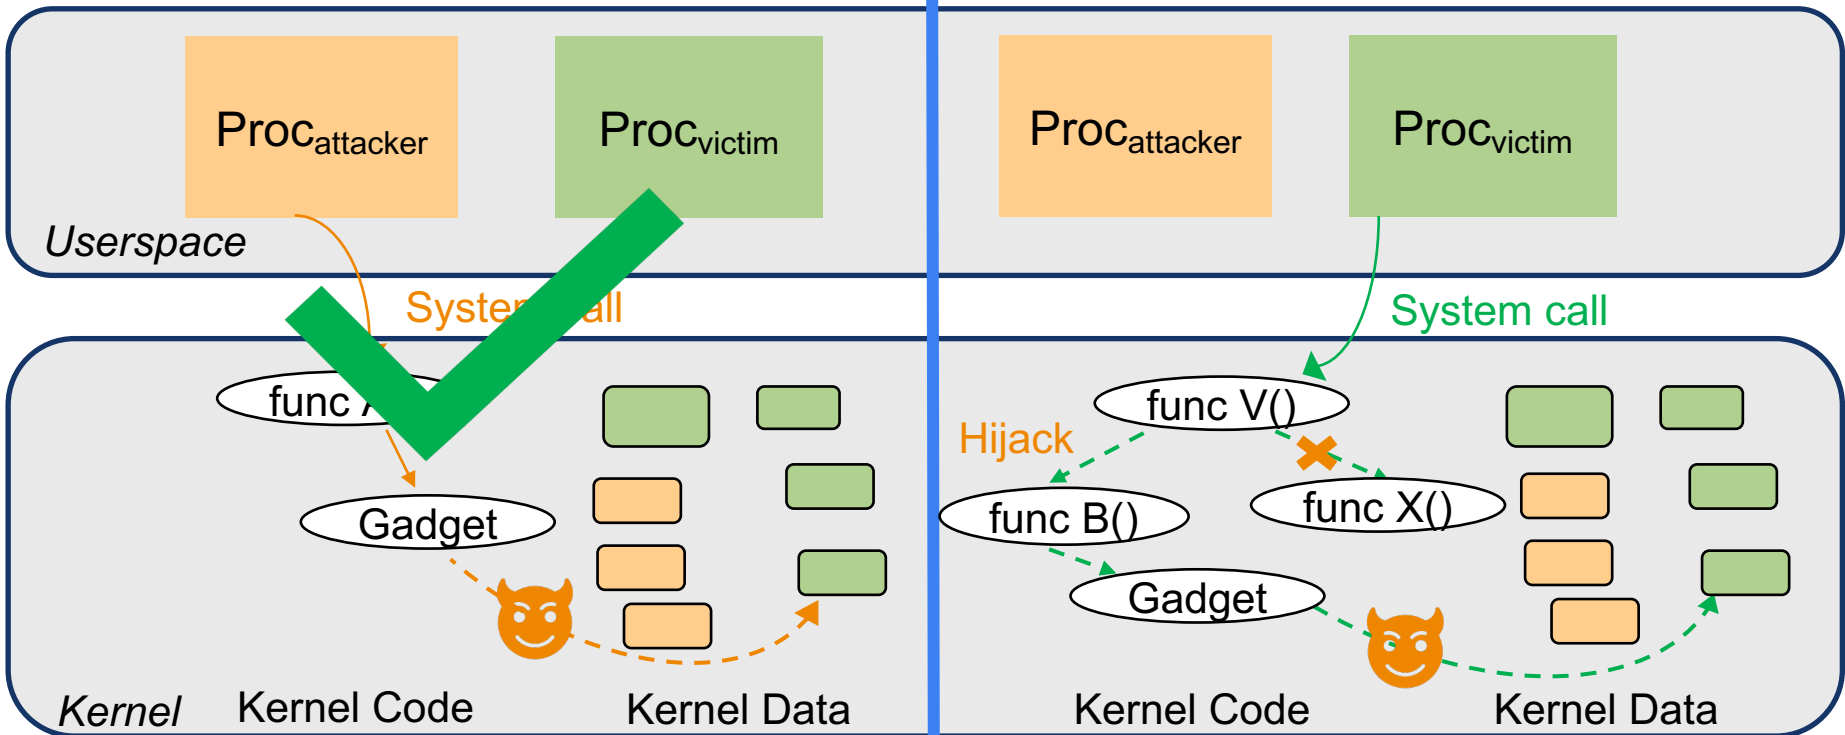


Perspective: Instruction Speculation Views

Transient ---→
 Non-transient —→

Active Attacks

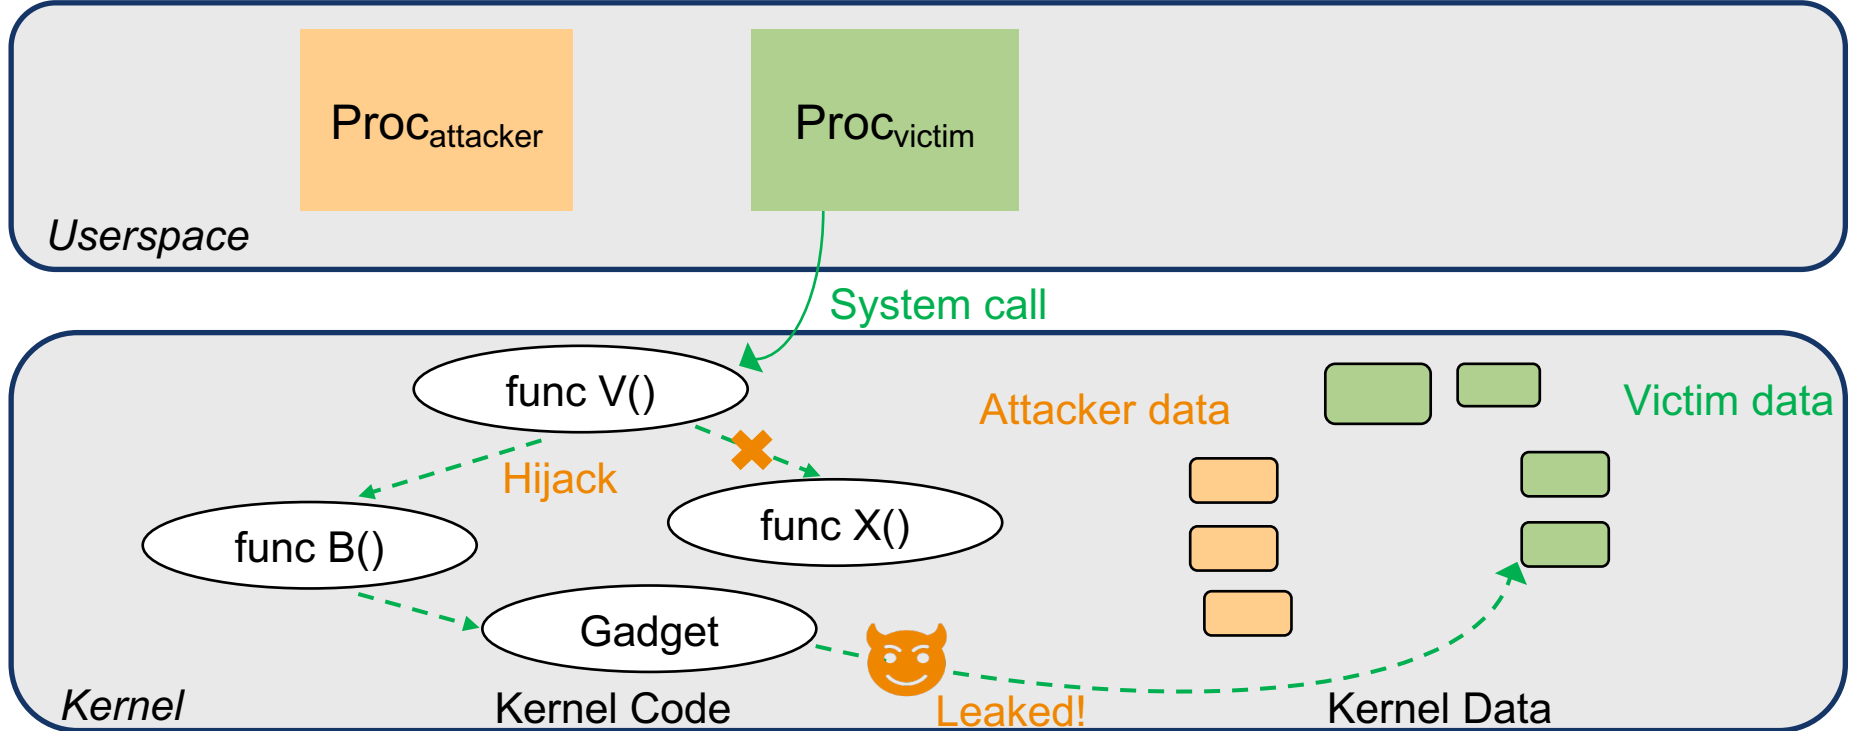
Passive Attacks



Perspective: Instruction Speculation Views

Transient 
Non-transient 

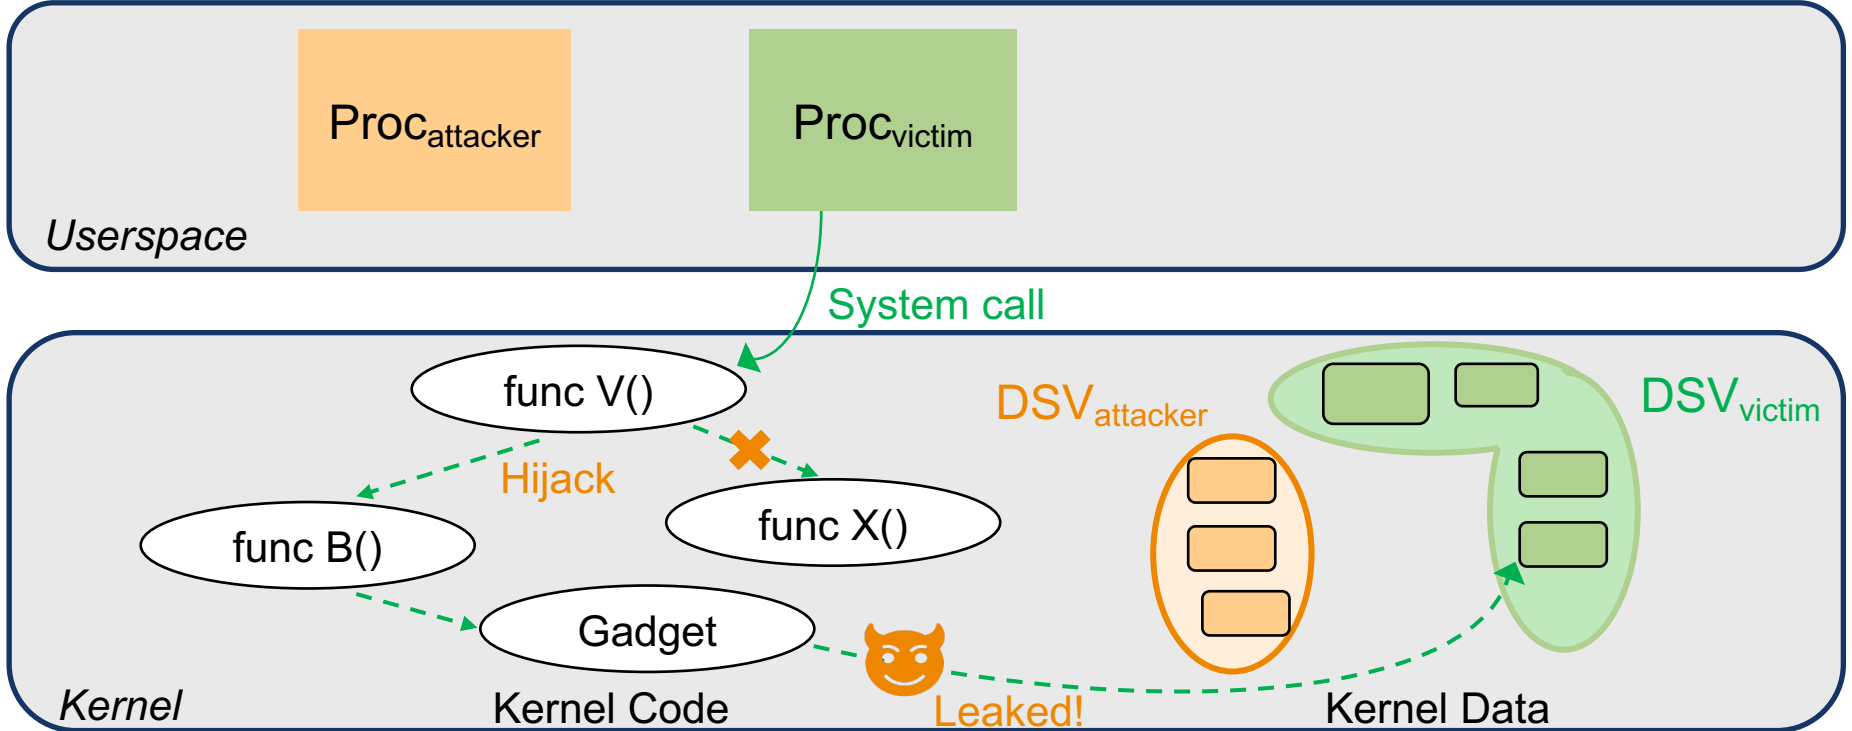
Passive Attacks



Perspective: Instruction Speculation Views

Transient 
Non-transient 

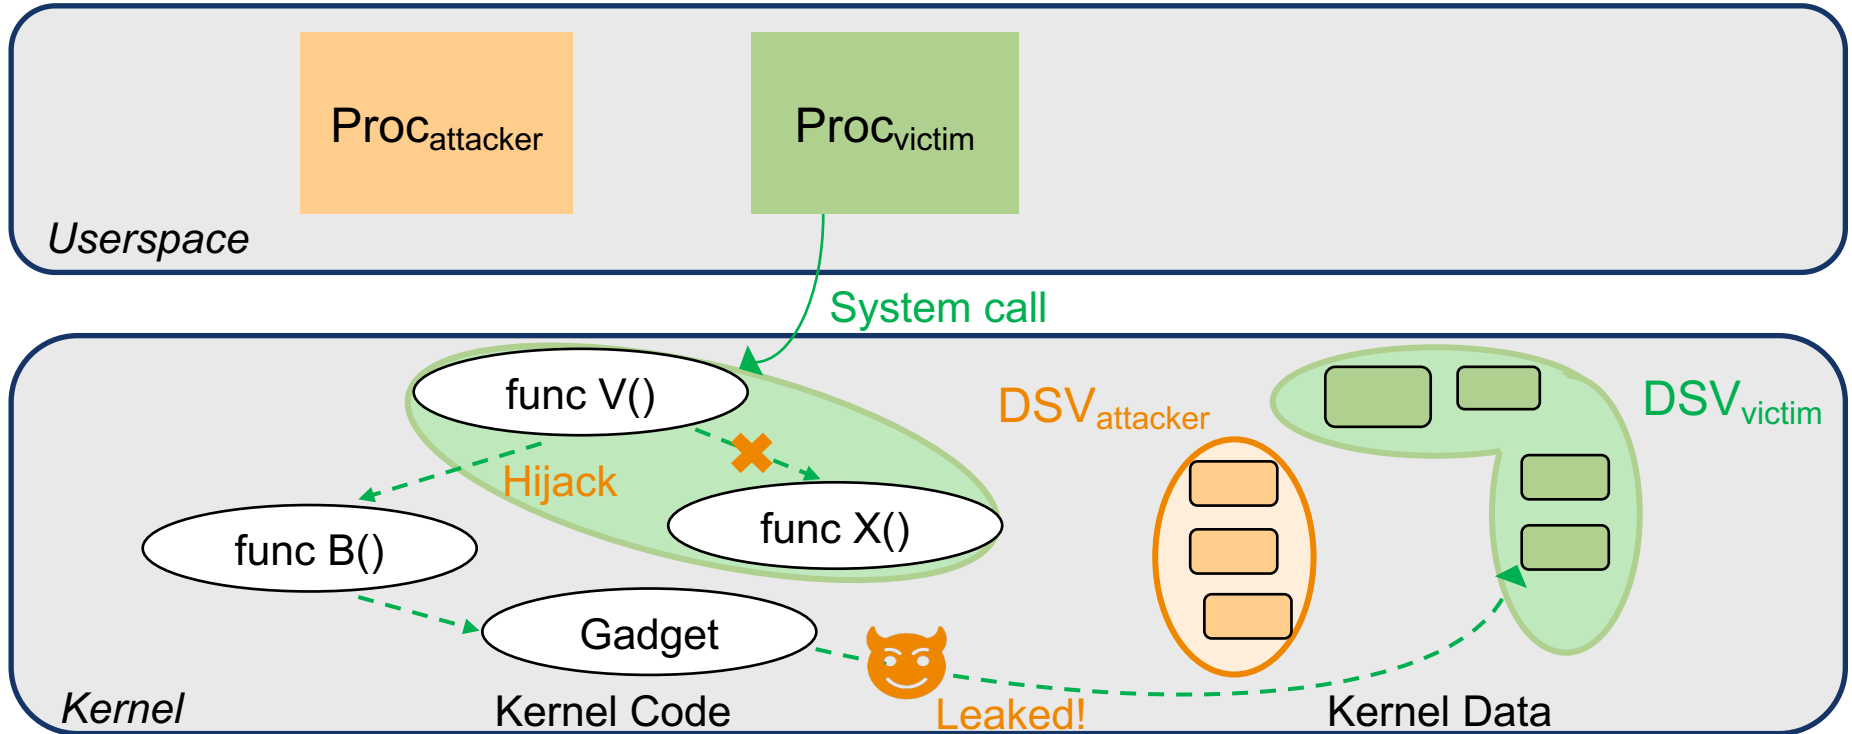
Passive Attacks



Perspective: Instruction Speculation Views

Transient 
Non-transient 

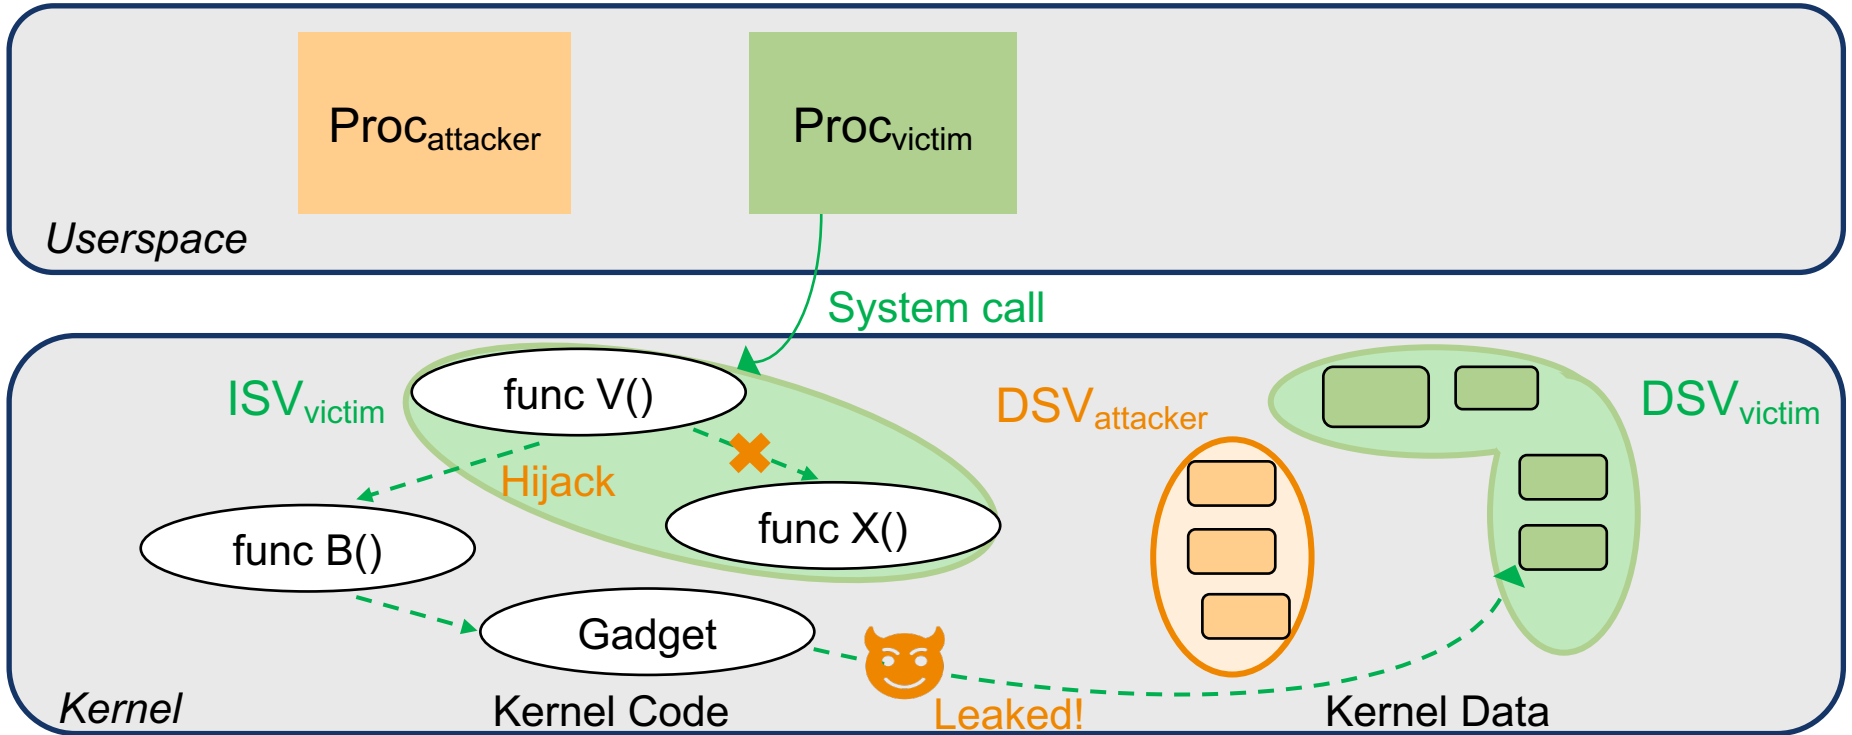
Passive Attacks



Perspective: Instruction Speculation Views

Transient 
Non-transient 

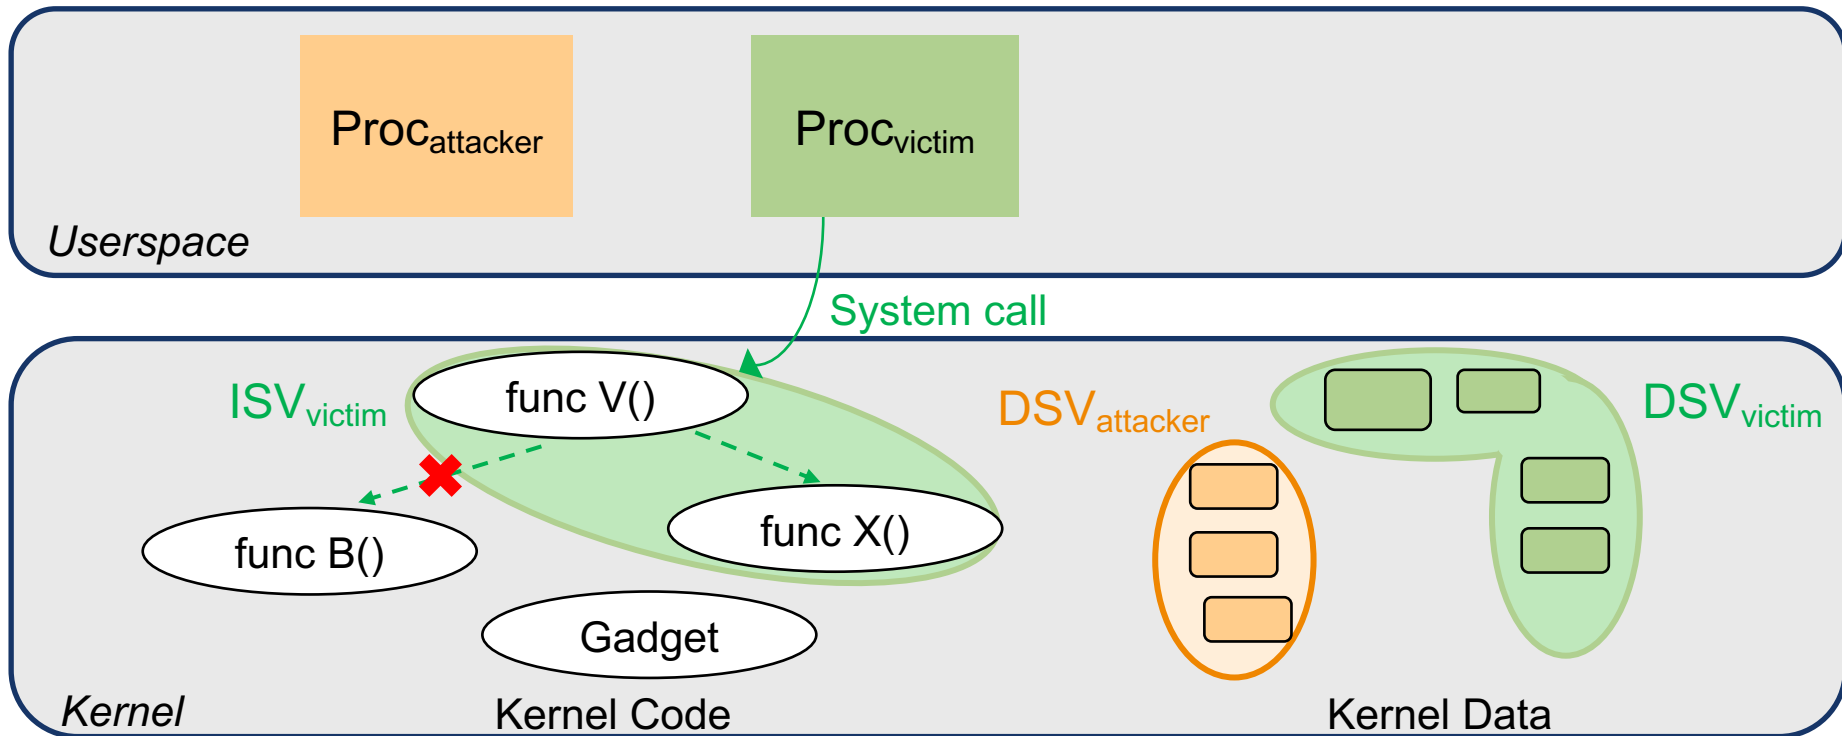
Passive Attacks



Perspective: Instruction Speculation Views

Transient ----→
Non-transient —→

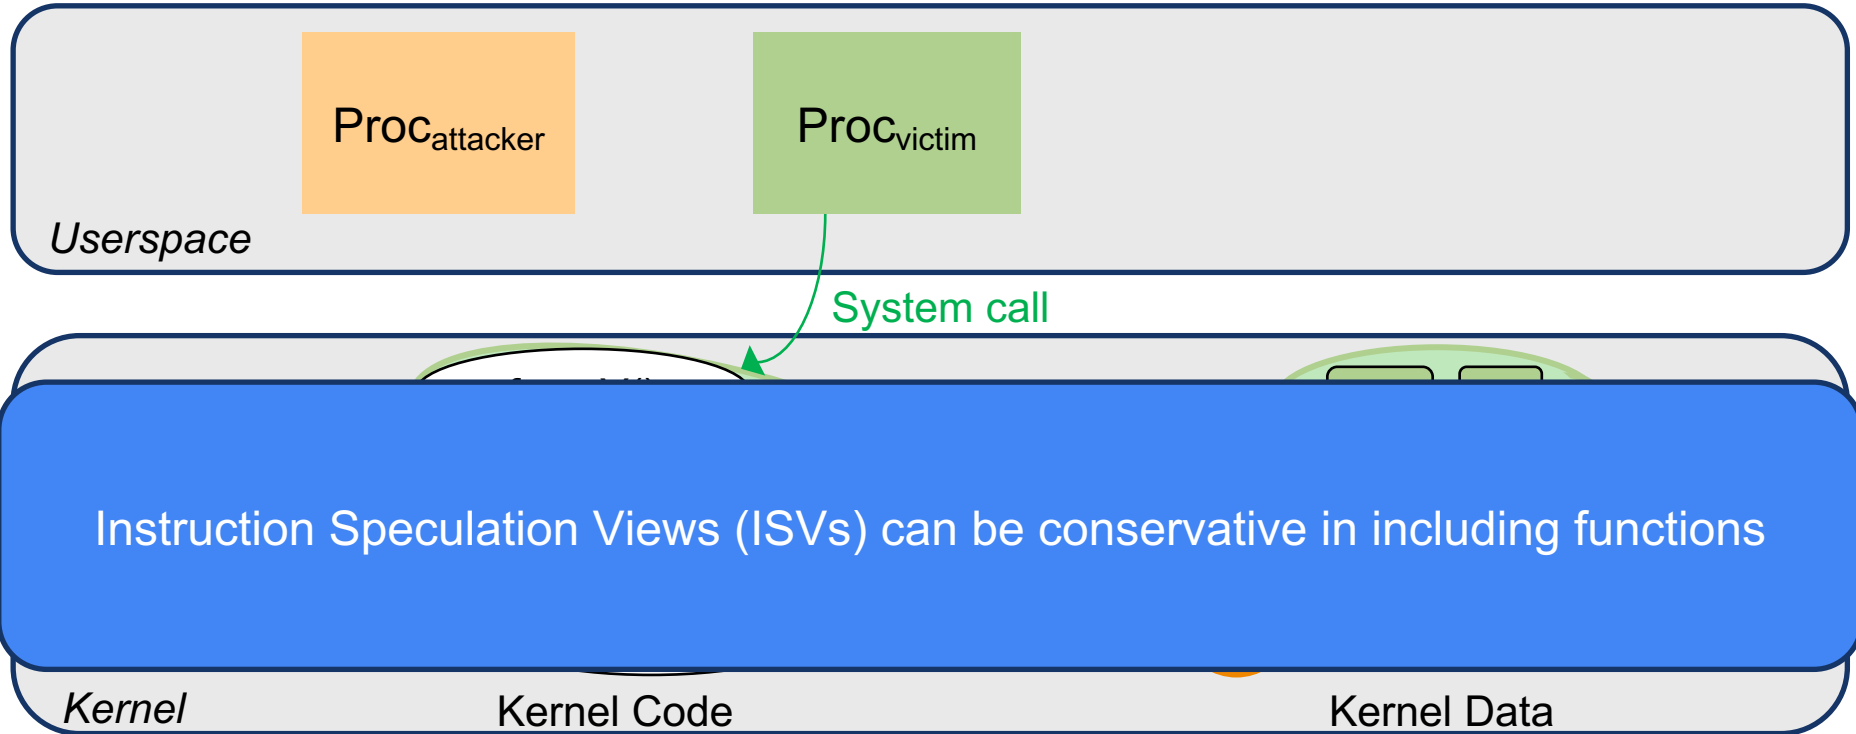
Passive Attacks



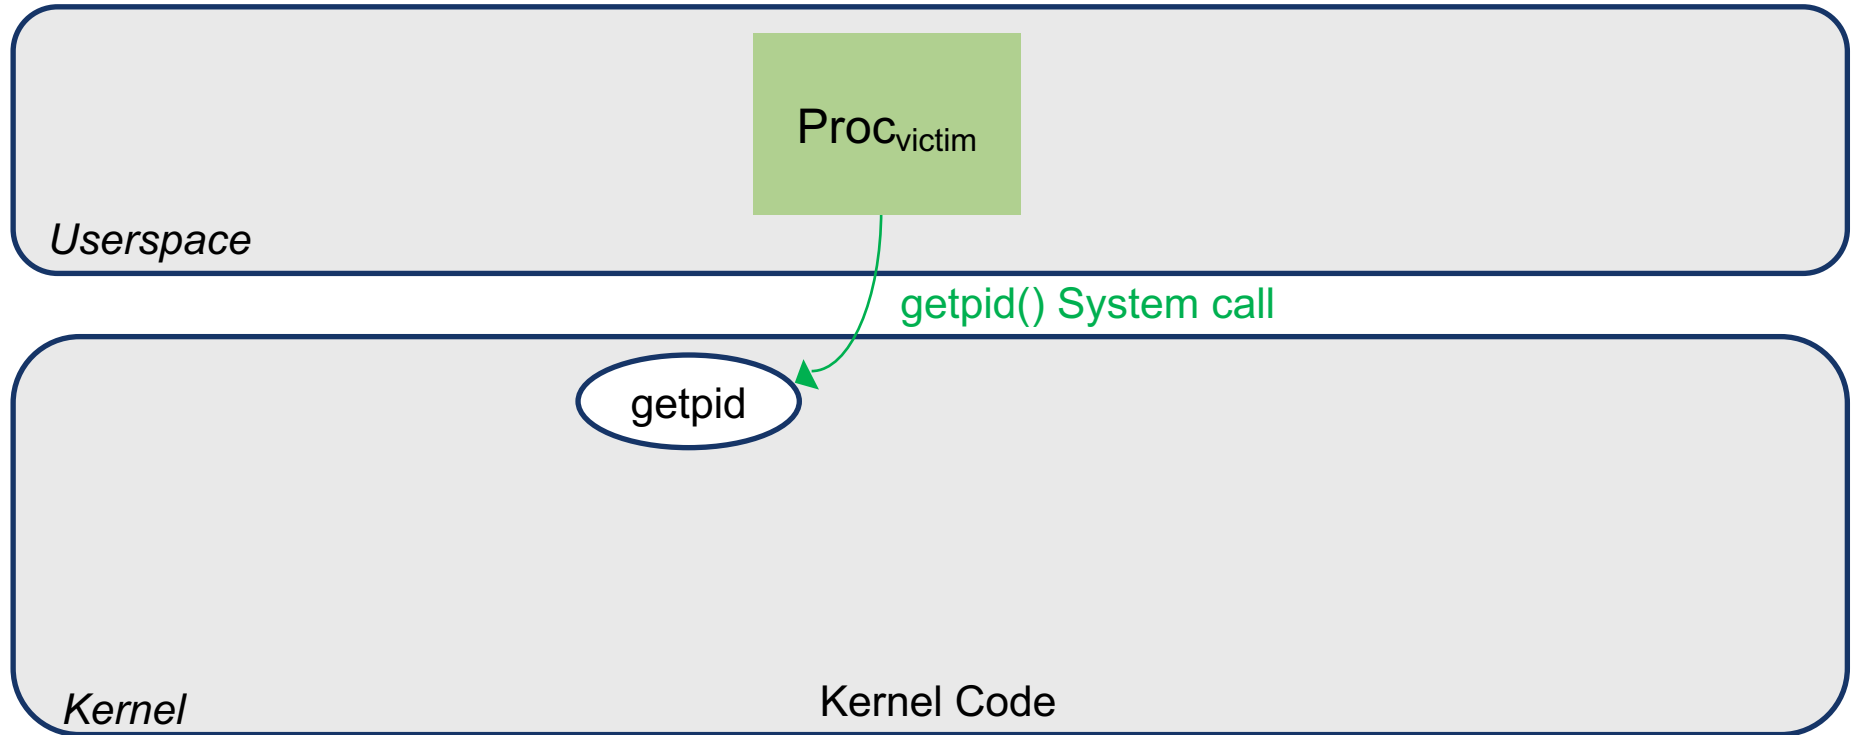
Perspective: Instruction Speculation Views

Transient 
Non-transient 

Passive Attacks

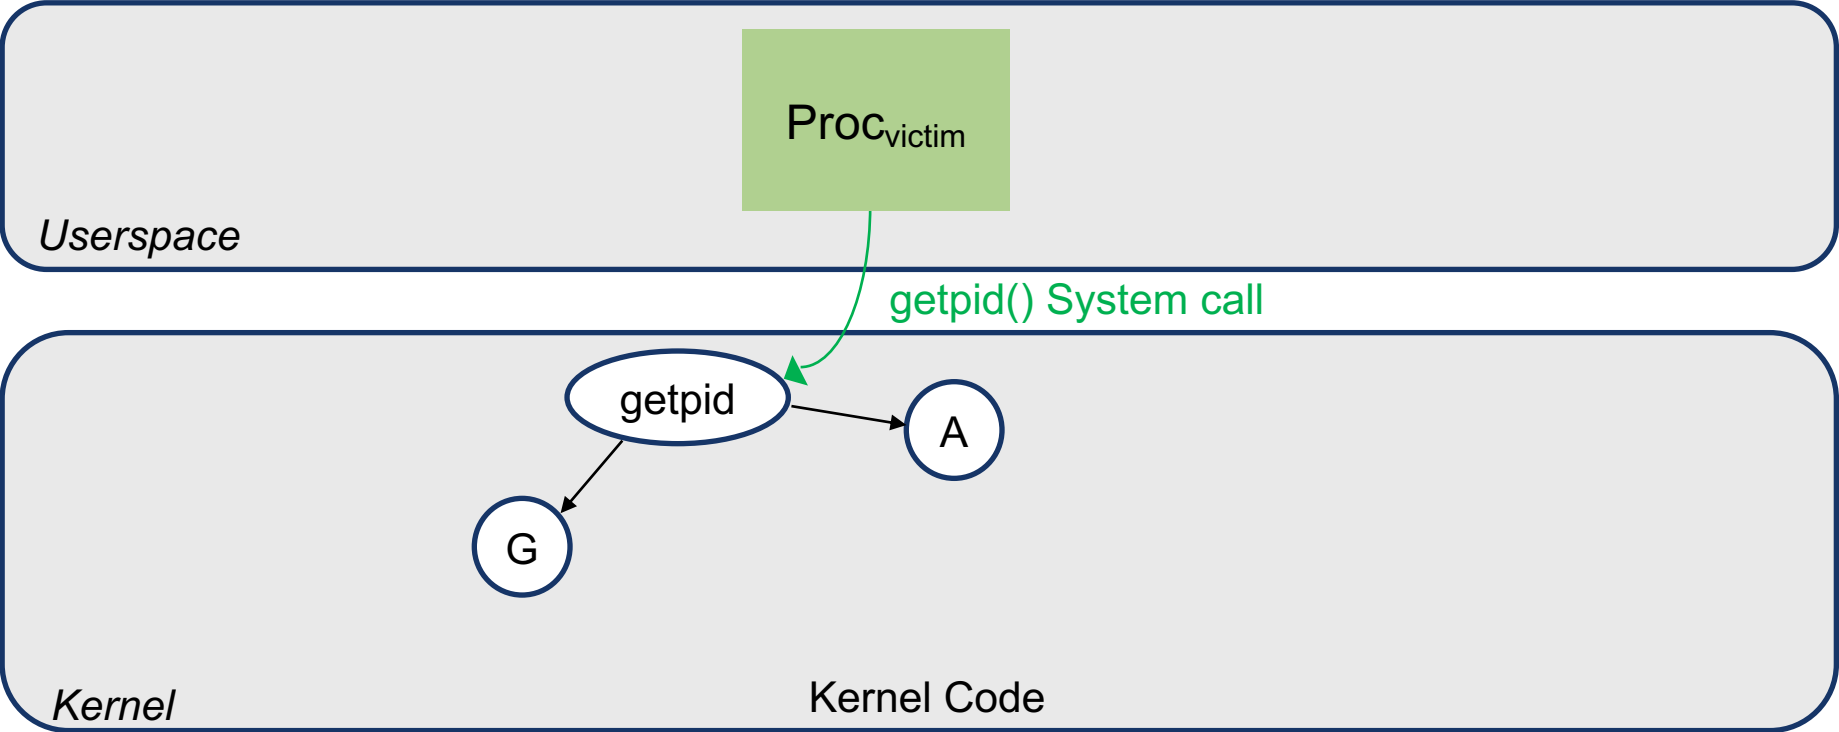


Generating ISV



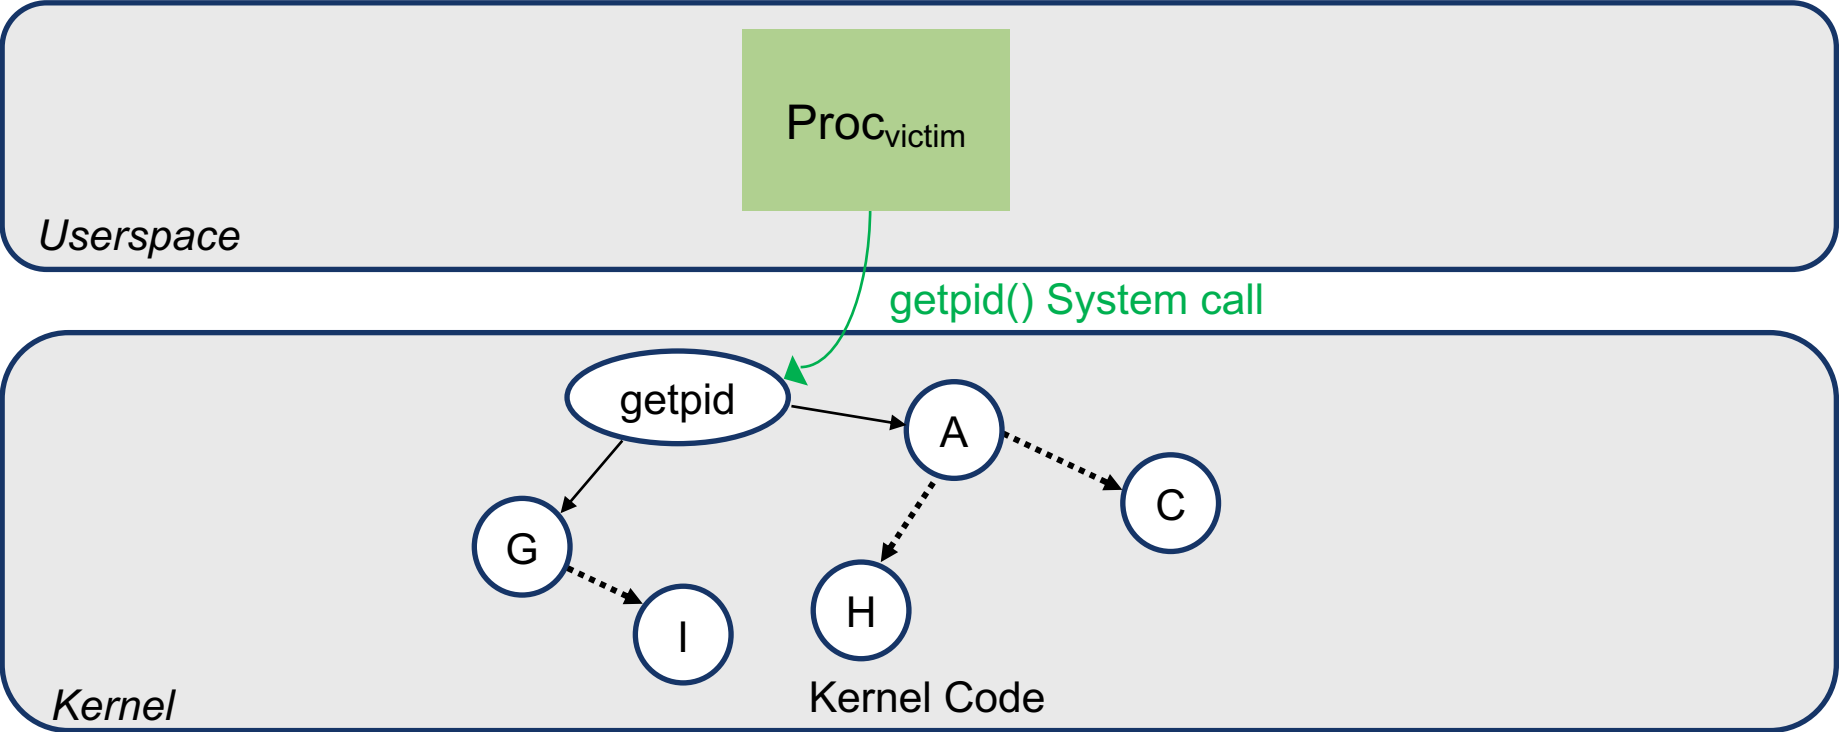
Generating ISV

Direct Call Edge \longrightarrow



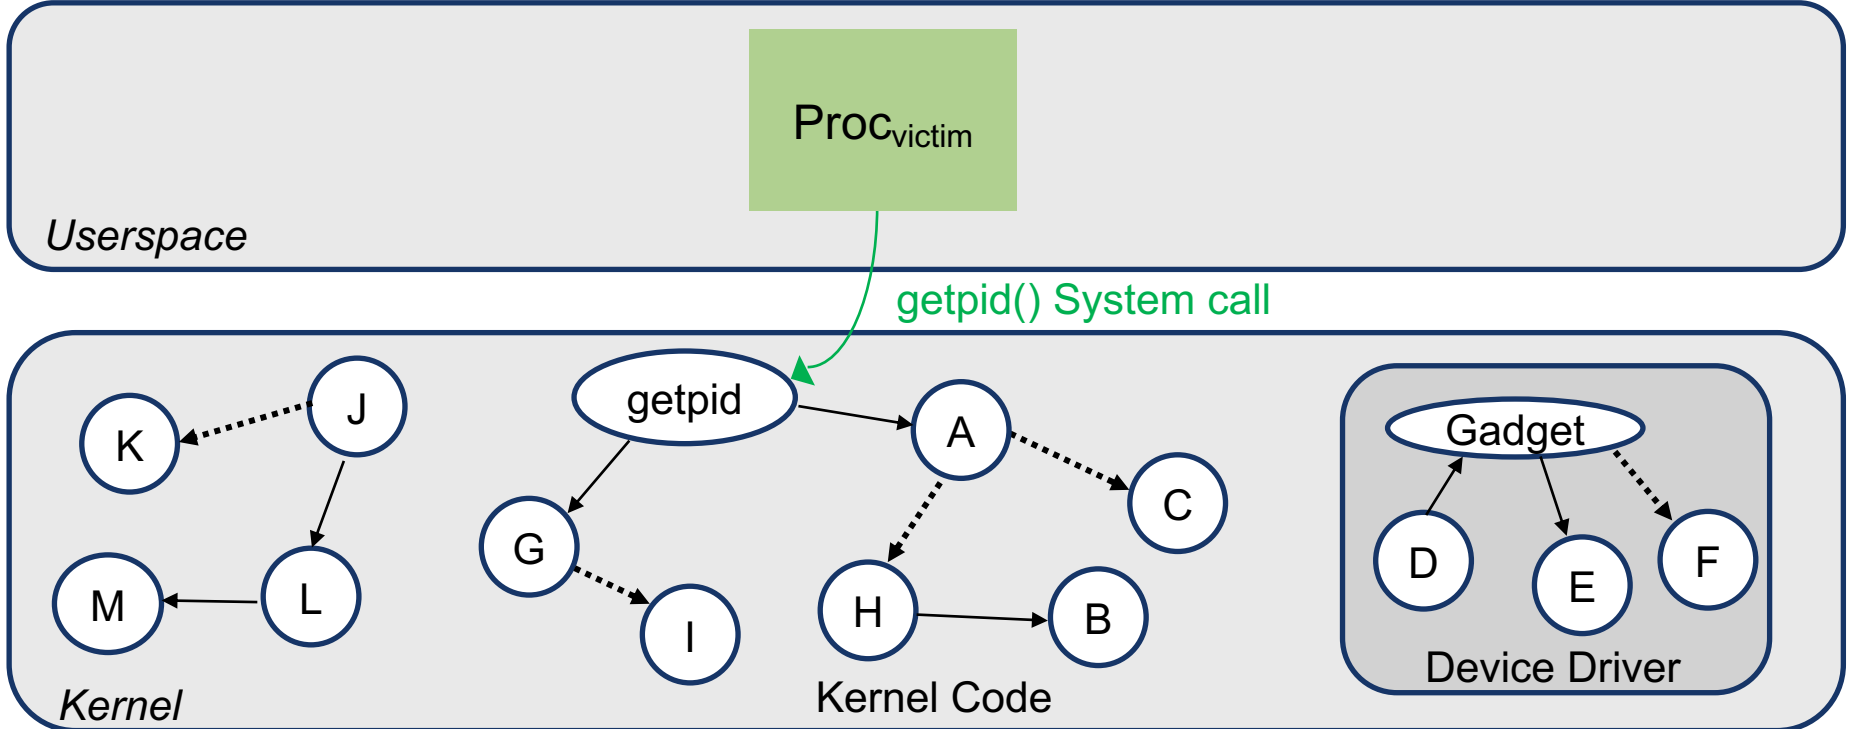
Generating ISV

Direct Call Edge \longrightarrow
Indirect Call Edge $\cdots\cdots\longrightarrow$



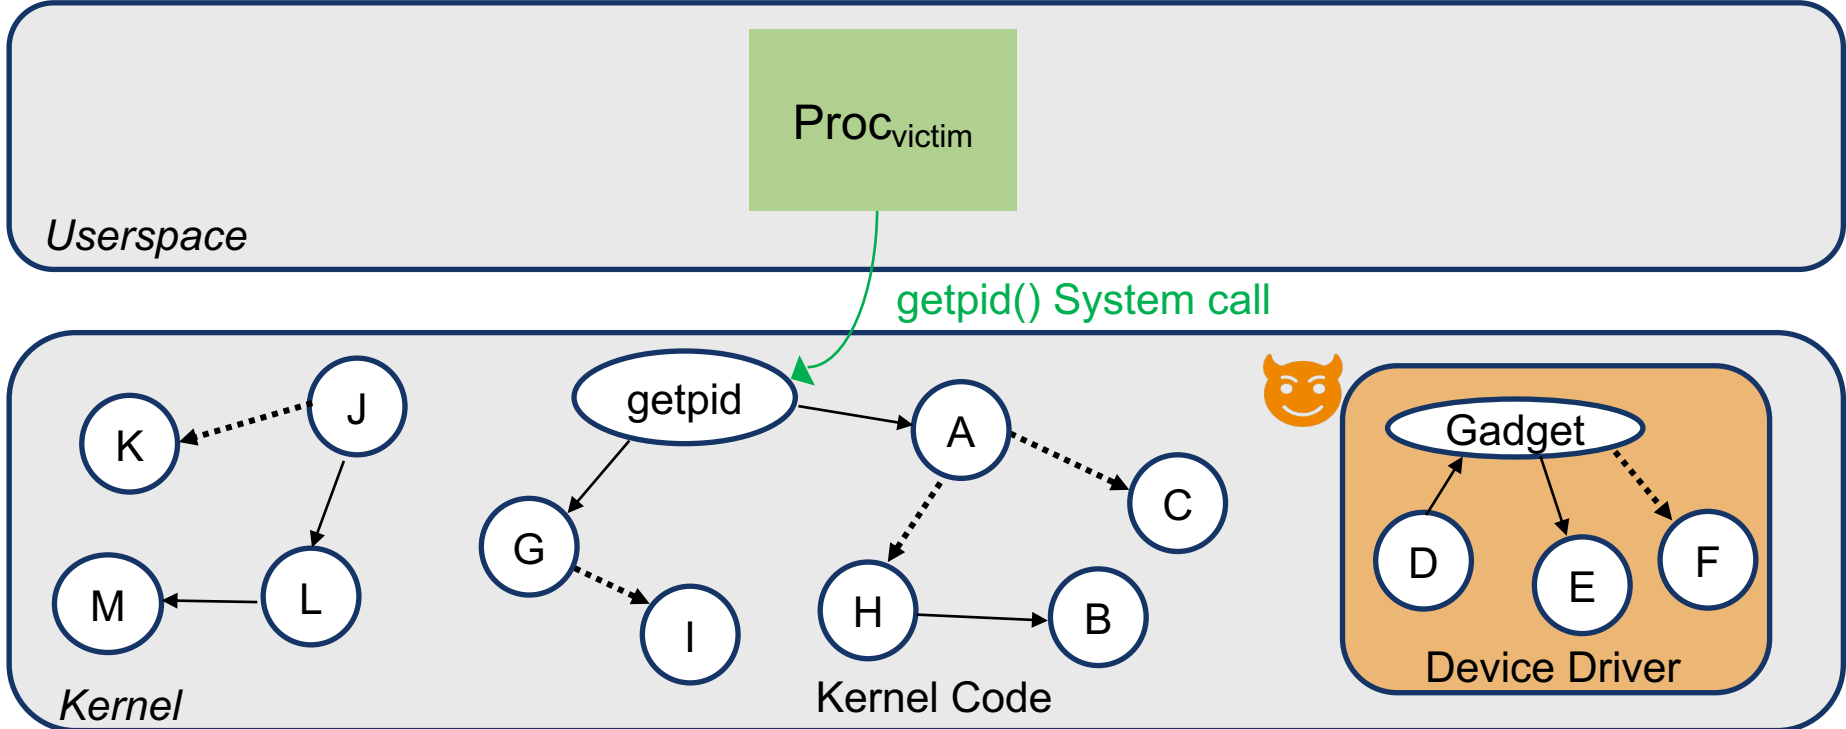
Generating ISV

Direct Call Edge \longrightarrow
Indirect Call Edge $\cdots\cdots\longrightarrow$



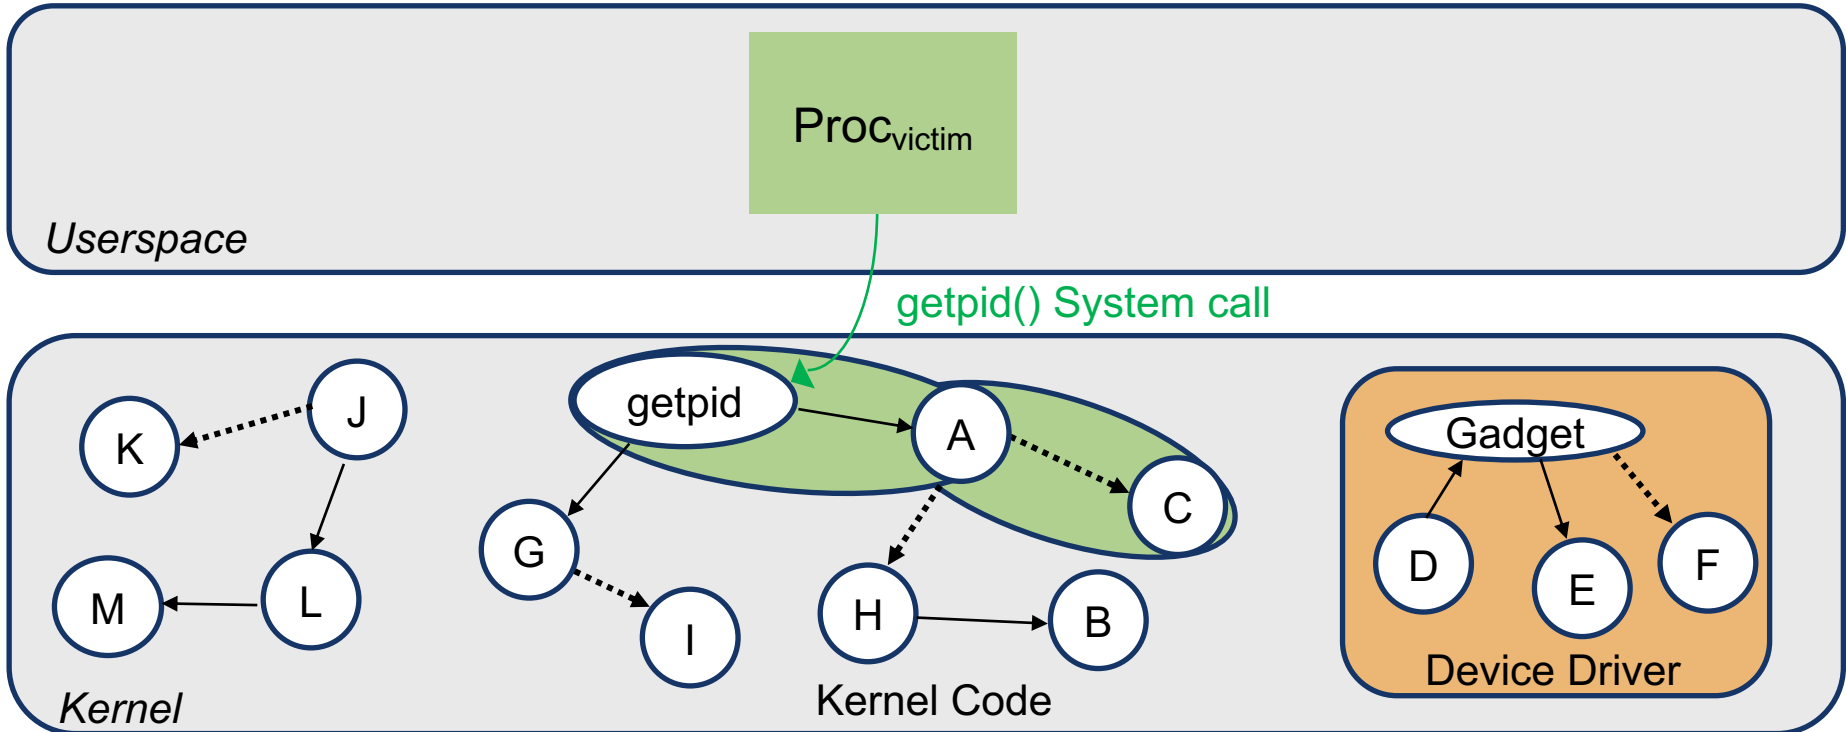
Generating ISV

Direct Call Edge \longrightarrow
Indirect Call Edge $\cdots\cdots\longrightarrow$



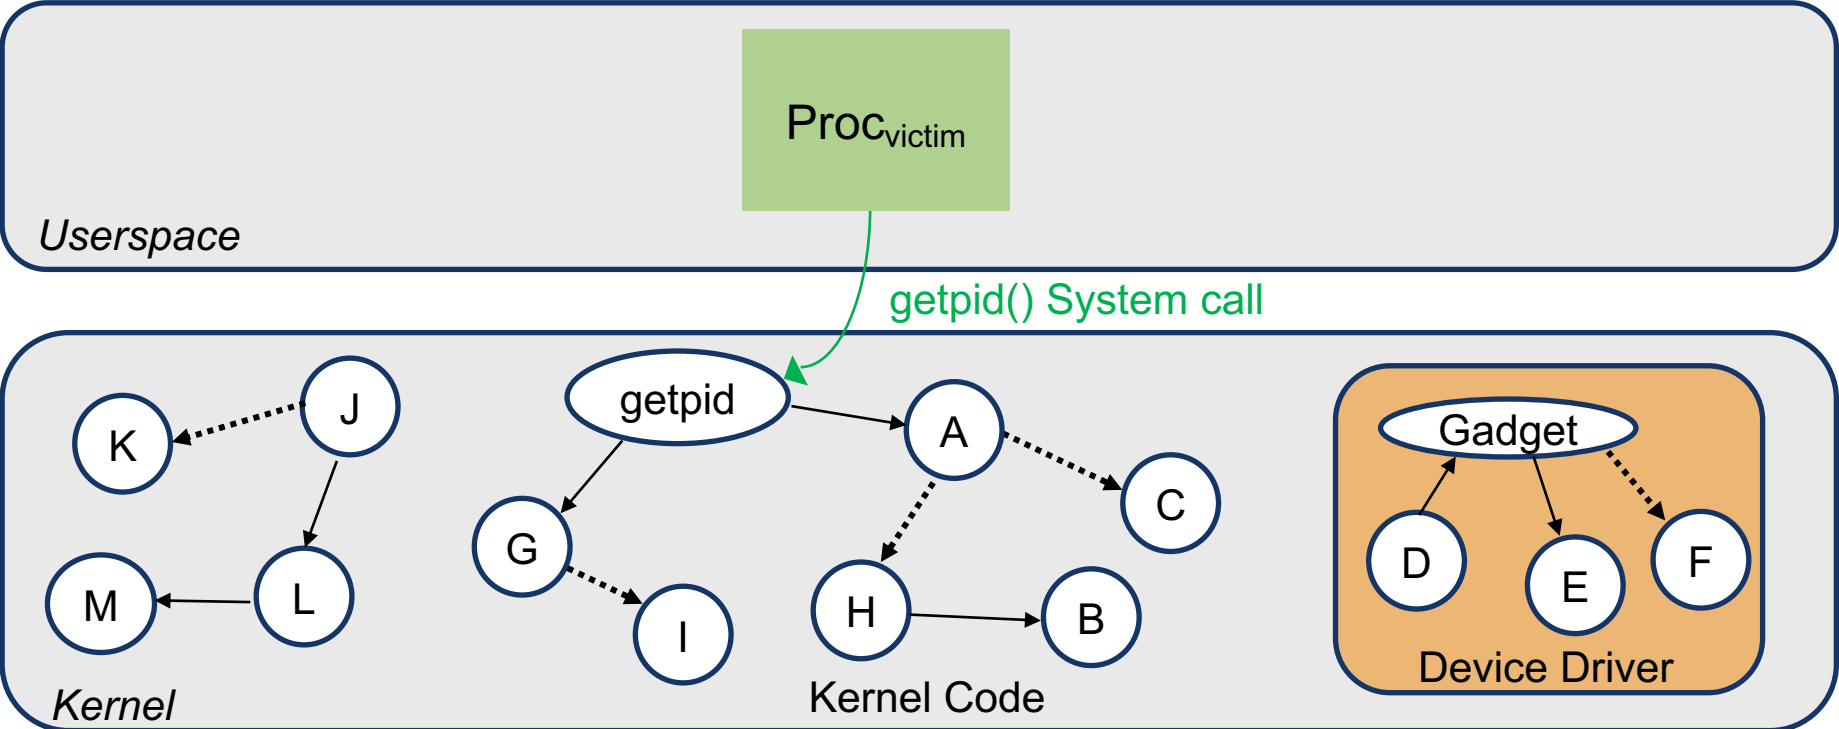
Generating ISV

Direct Call Edge \longrightarrow
Indirect Call Edge $\cdots\cdots\longrightarrow$



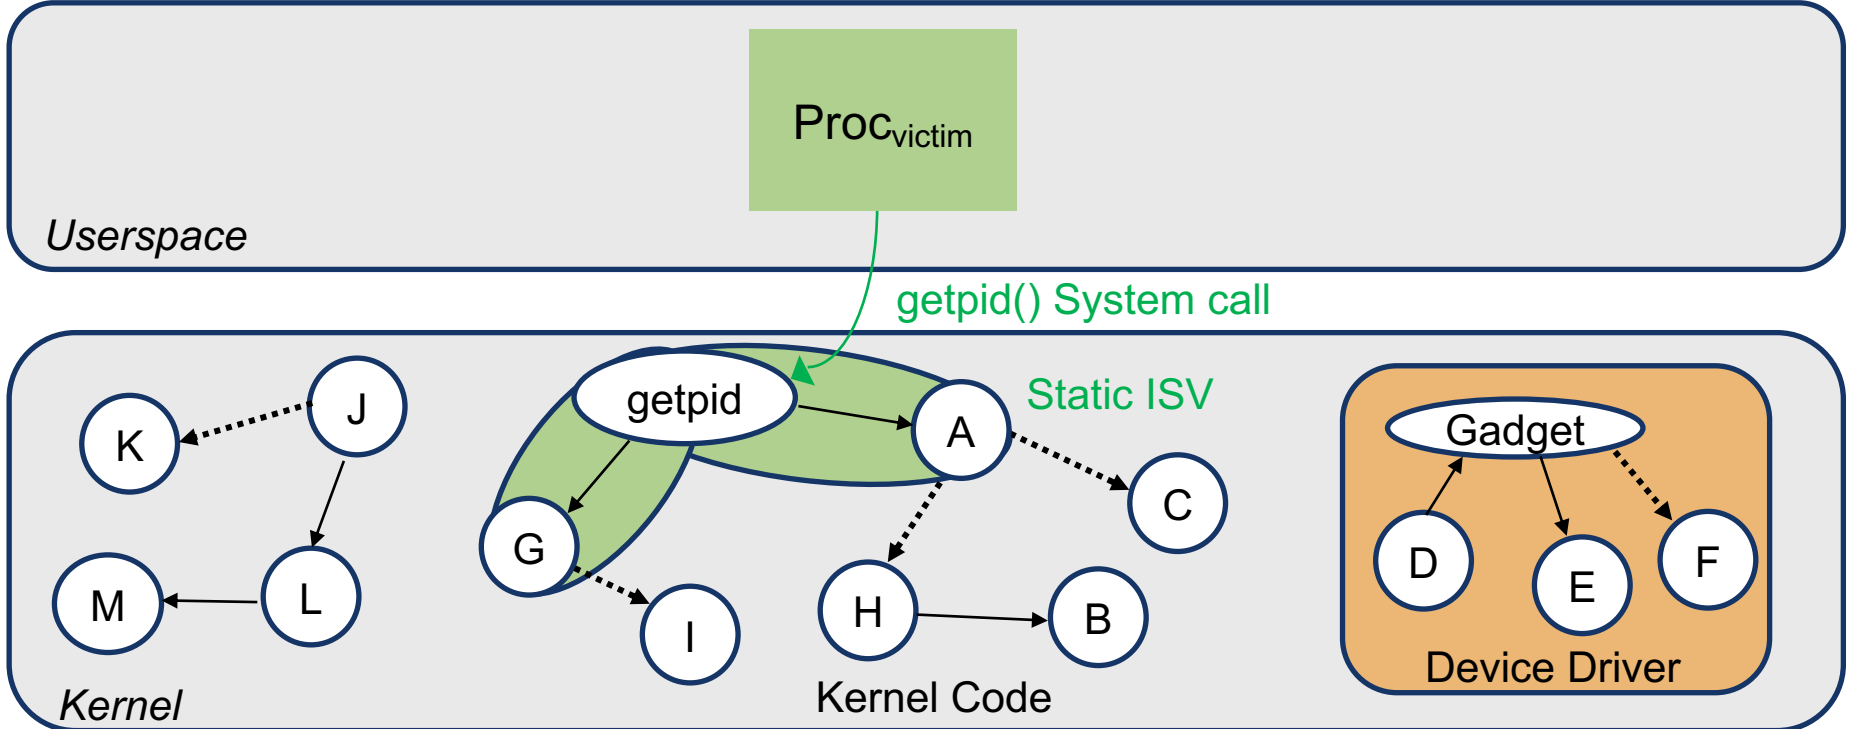
Static ISV

Direct Call Edge \longrightarrow
Indirect Call Edge $\cdots\cdots\longrightarrow$



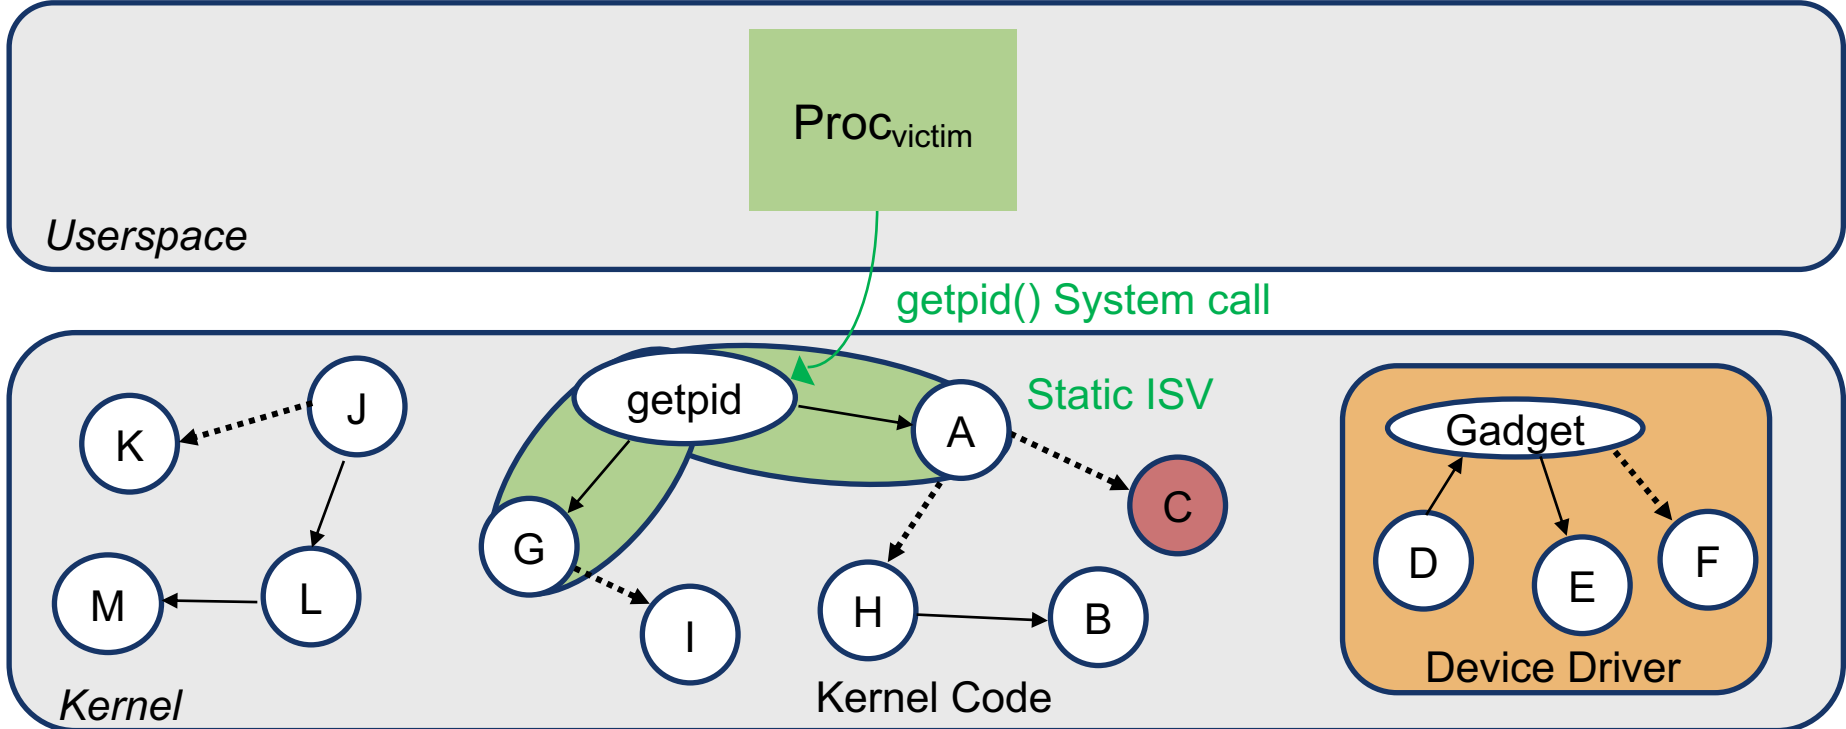
Static ISV

Direct Call Edge \longrightarrow
Indirect Call Edge $\cdots\cdots\longrightarrow$



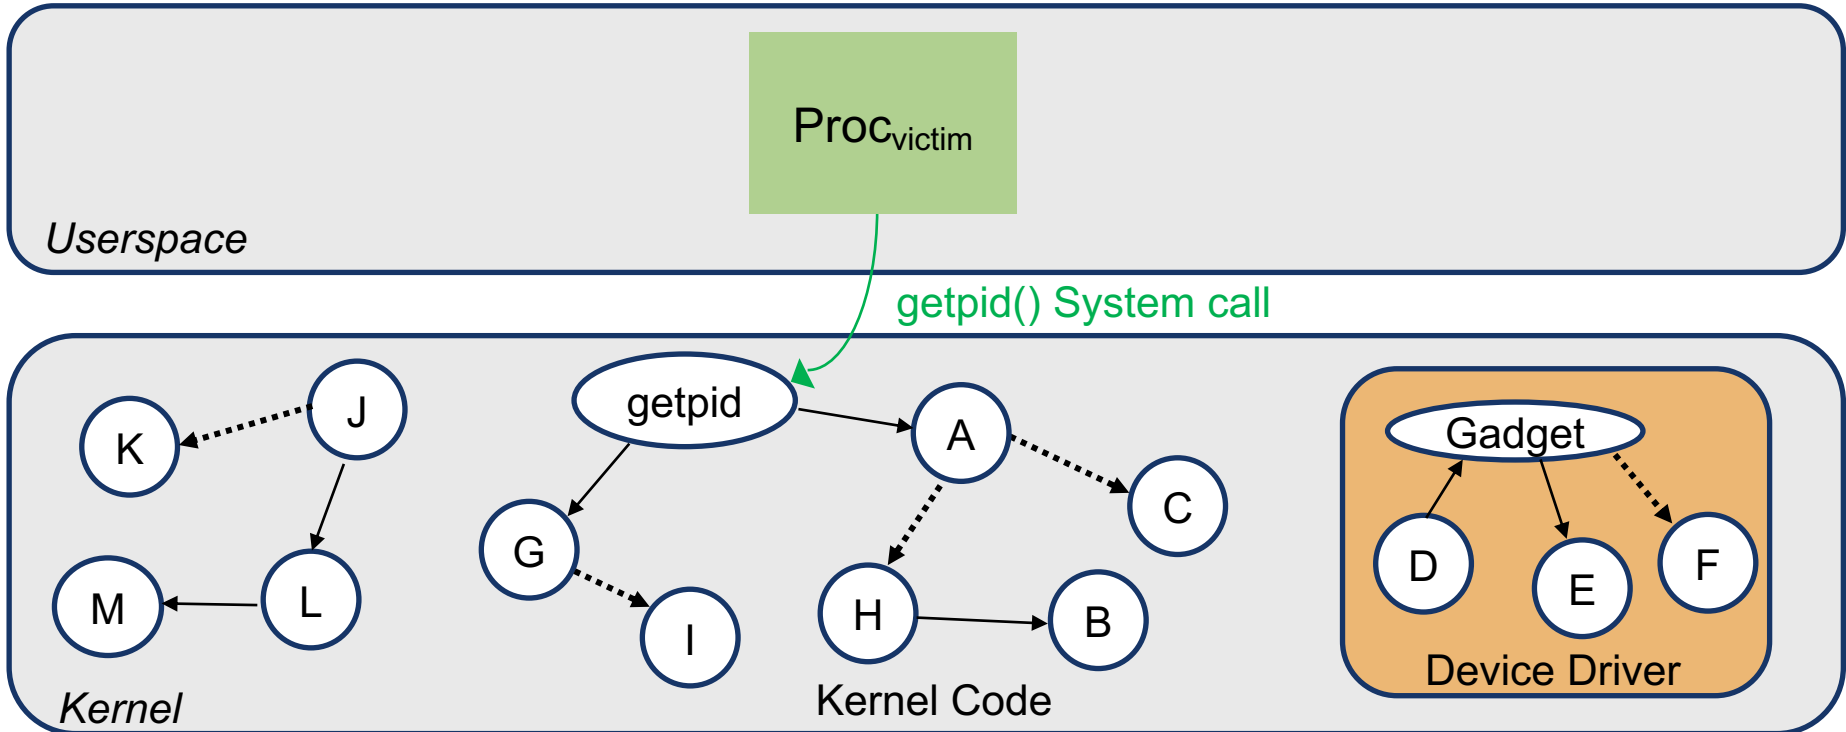
Static ISV

Direct Call Edge \longrightarrow
Indirect Call Edge $\cdots\cdots\longrightarrow$



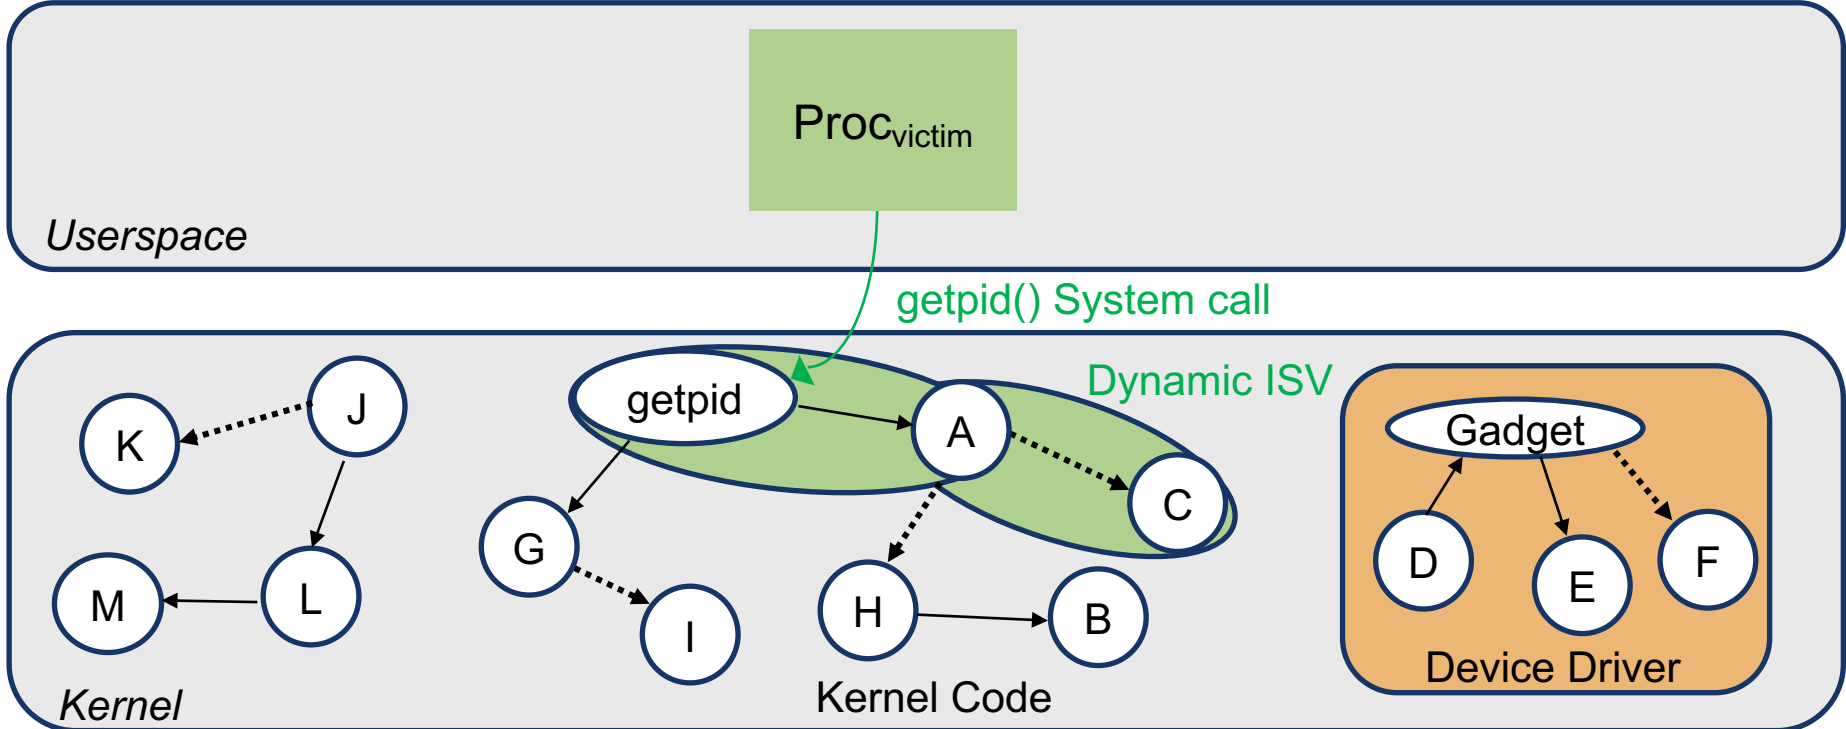
Dynamic ISV

Direct Call Edge \longrightarrow
Indirect Call Edge $\cdots\cdots\longrightarrow$



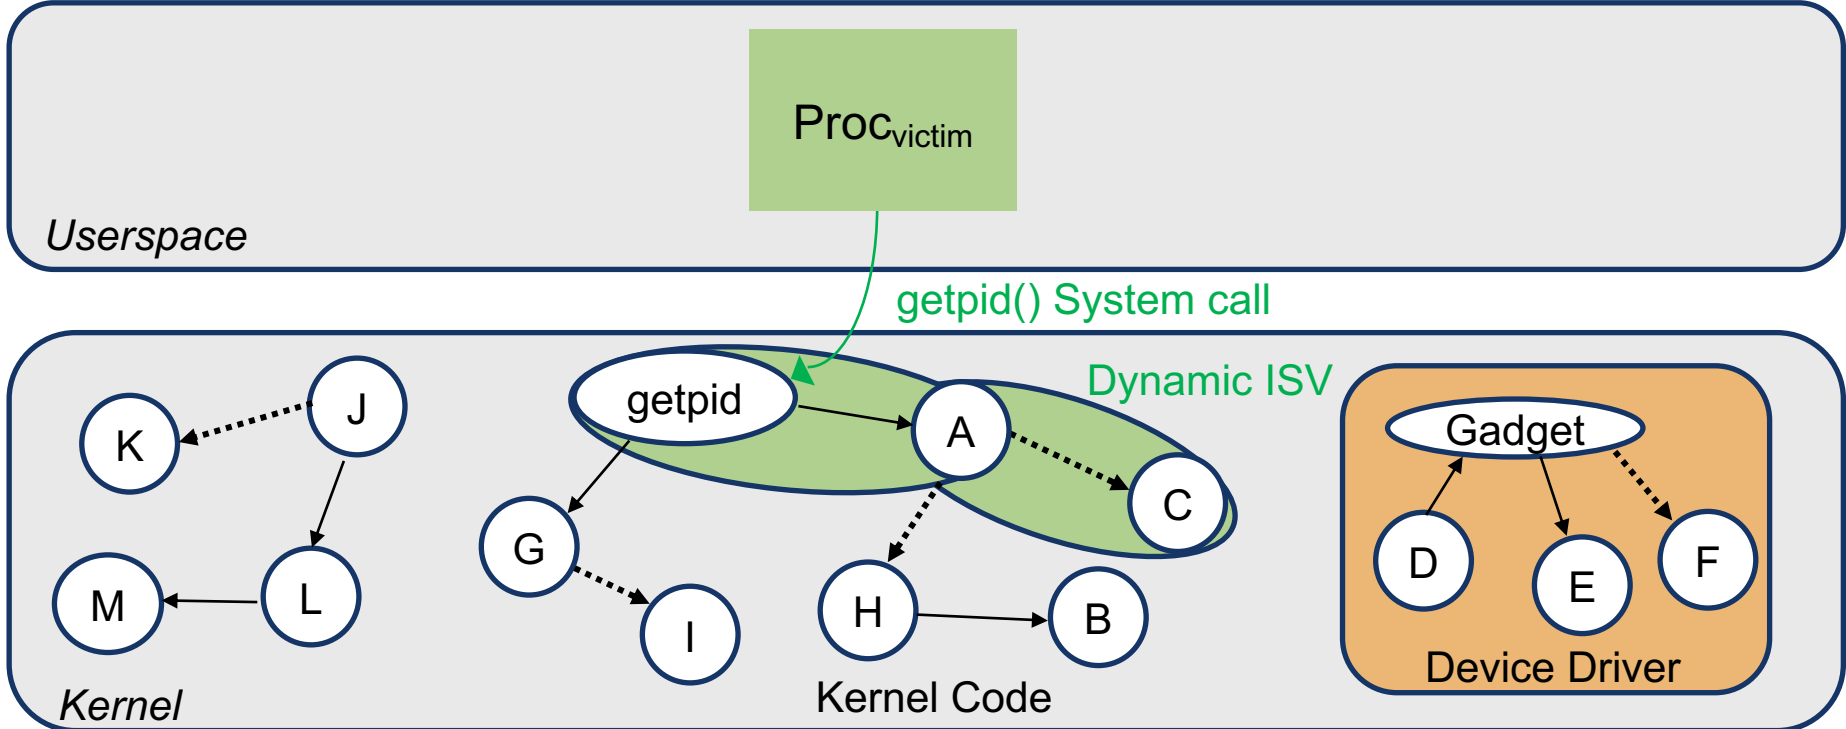
Dynamic ISV

Direct Call Edge \longrightarrow
Indirect Call Edge $\cdots\cdots\longrightarrow$



Benefits of ISVs

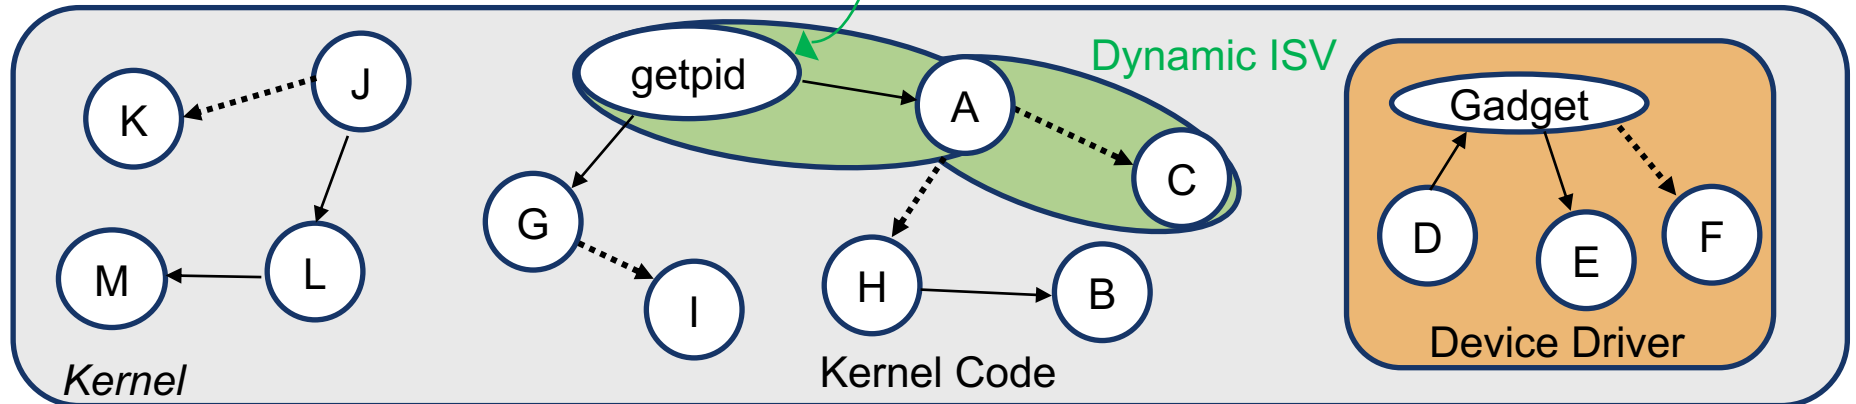
Direct Call Edge \longrightarrow
Indirect Call Edge $\cdots\cdots\longrightarrow$



Benefits of ISVs

Direct Call Edge \longrightarrow
Indirect Call Edge $\cdots\cdots\cdots\longrightarrow$

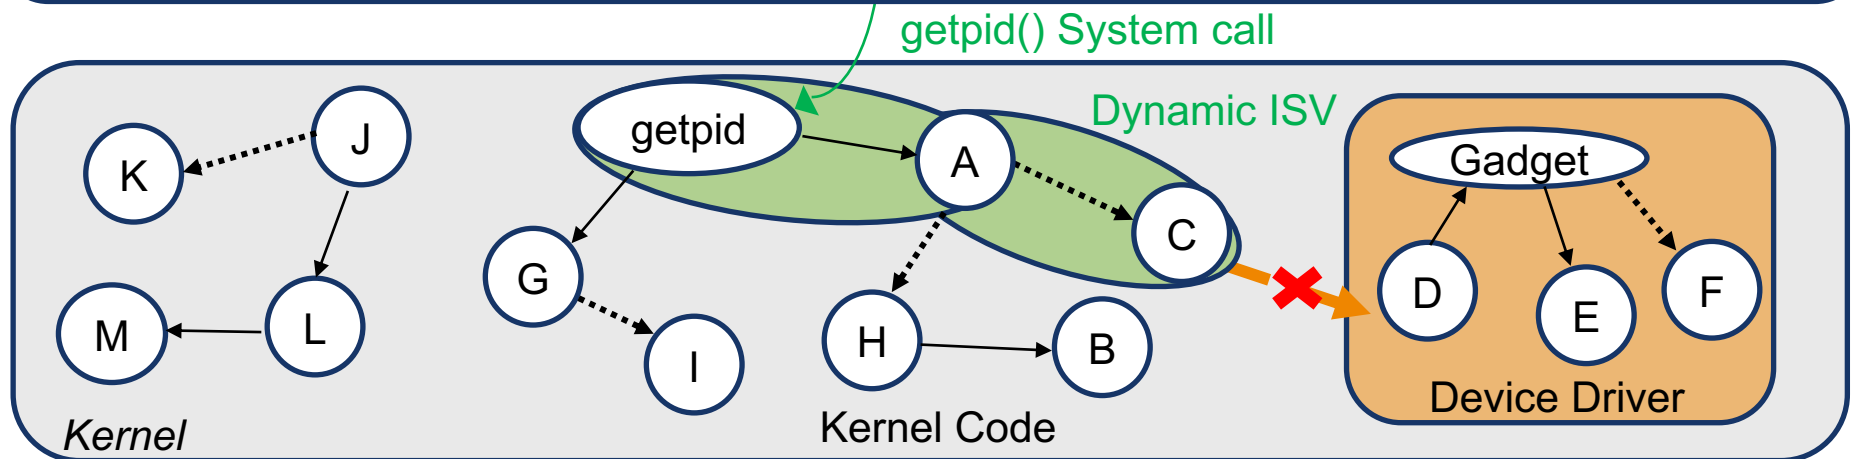
Reduce Attack Surface: Exclude unused or infrequently used kernel functions



Benefits of ISVs

Direct Call Edge \longrightarrow
Indirect Call Edge $\cdots\cdots\cdots\longrightarrow$

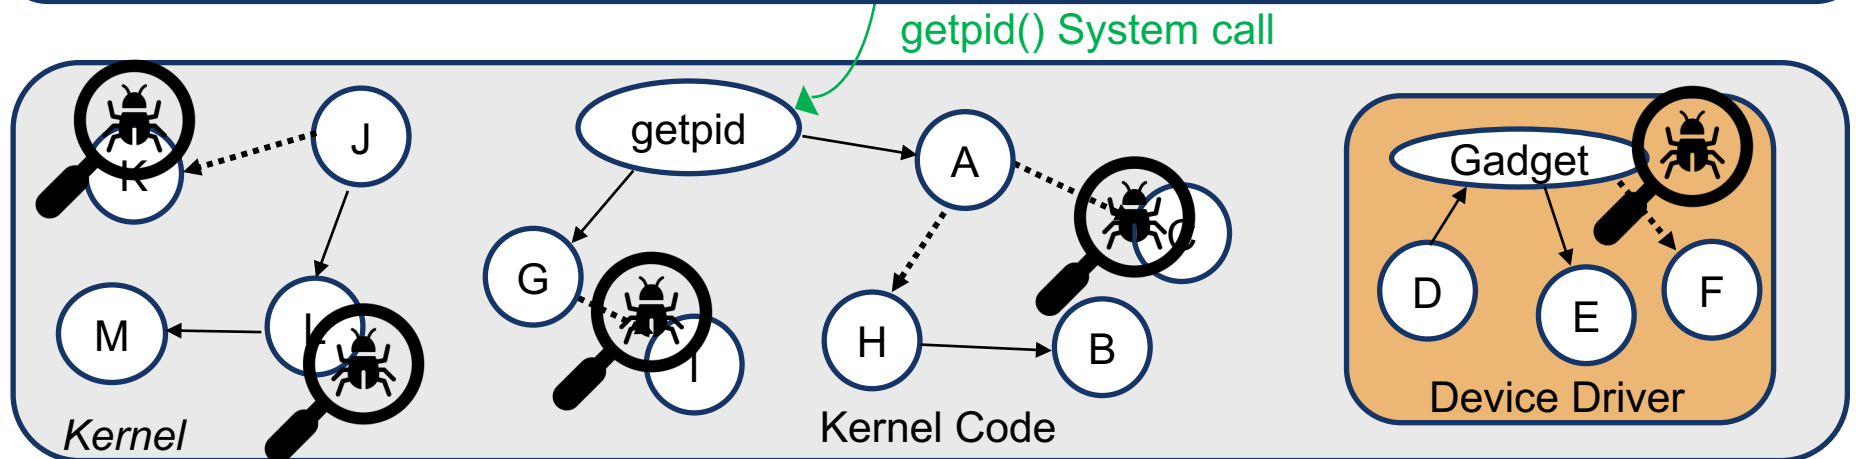
Reduce Attack Surface: Exclude unused or infrequently used kernel functions



Benefits of ISVs

Direct Call Edge \longrightarrow
Indirect Call Edge $\cdots\cdots\longrightarrow$

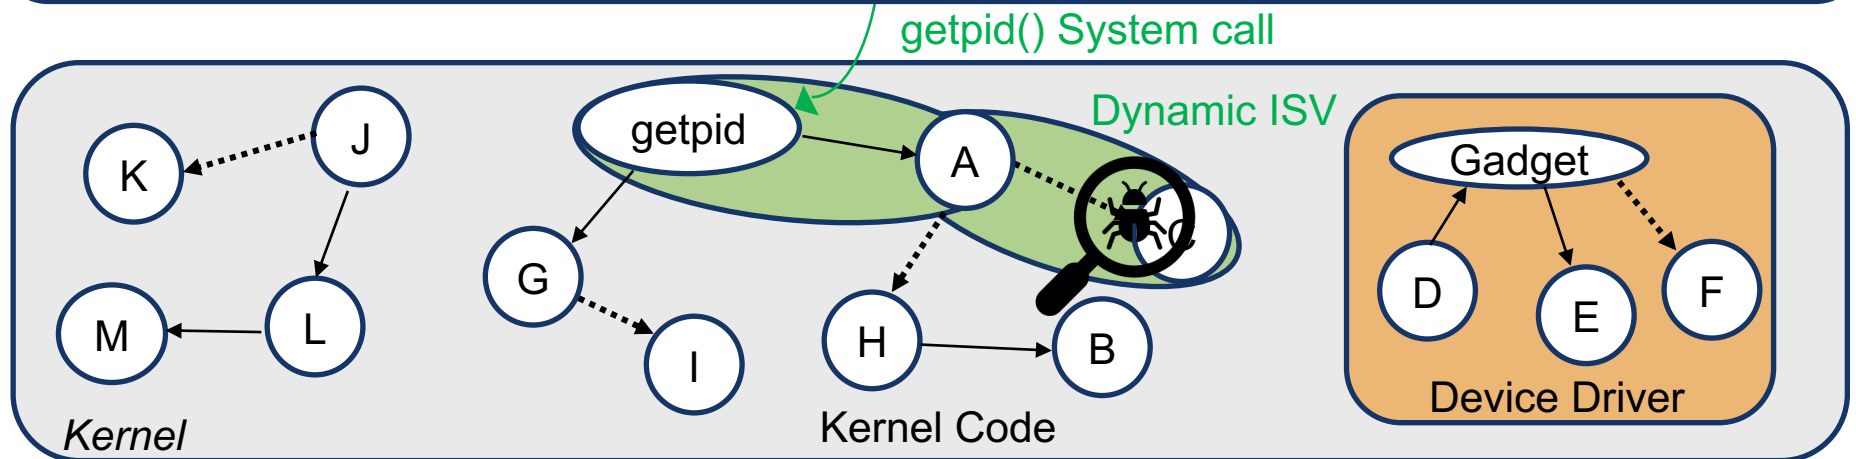
Accelerate Gadget Finding: Focus on a smaller set of functions



Benefits of ISVs

Direct Call Edge \longrightarrow
Indirect Call Edge $\cdots\cdots\cdots\longrightarrow$

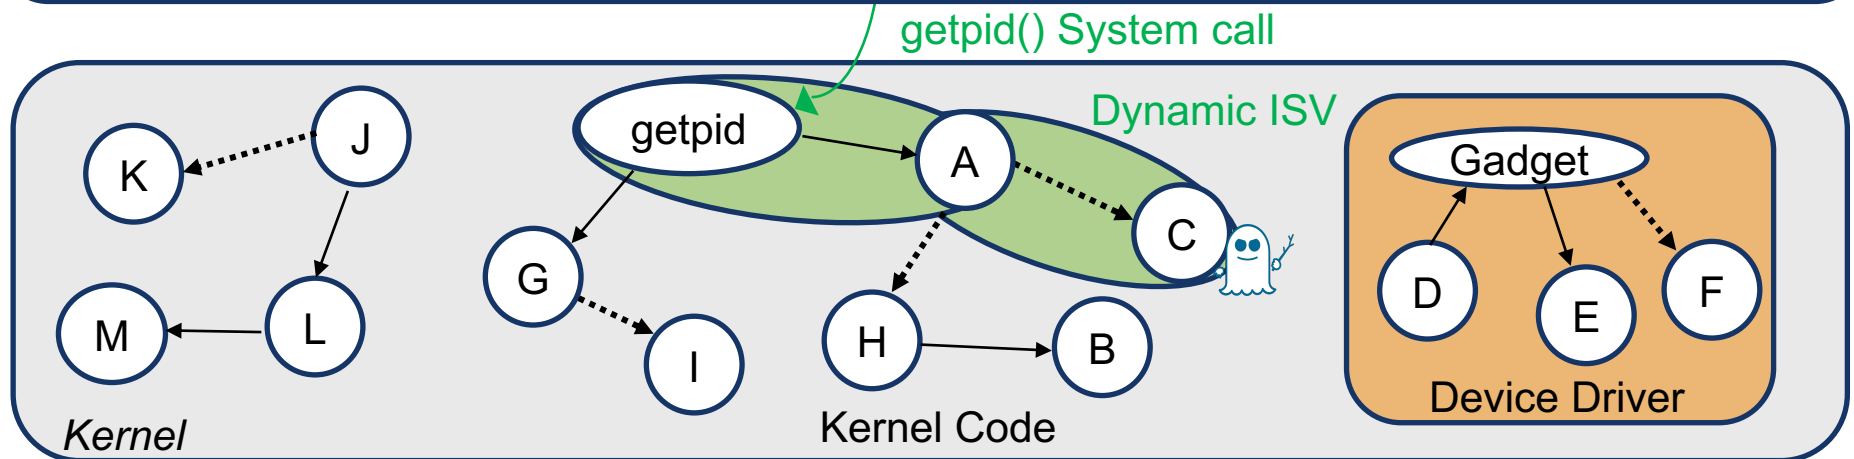
Accelerate Gadget Finding: Focus on a smaller set of functions



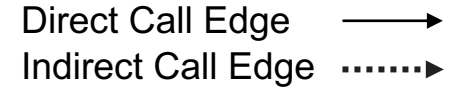
Benefits of ISVs

Direct Call Edge \longrightarrow
Indirect Call Edge $\cdots\cdots\longrightarrow$

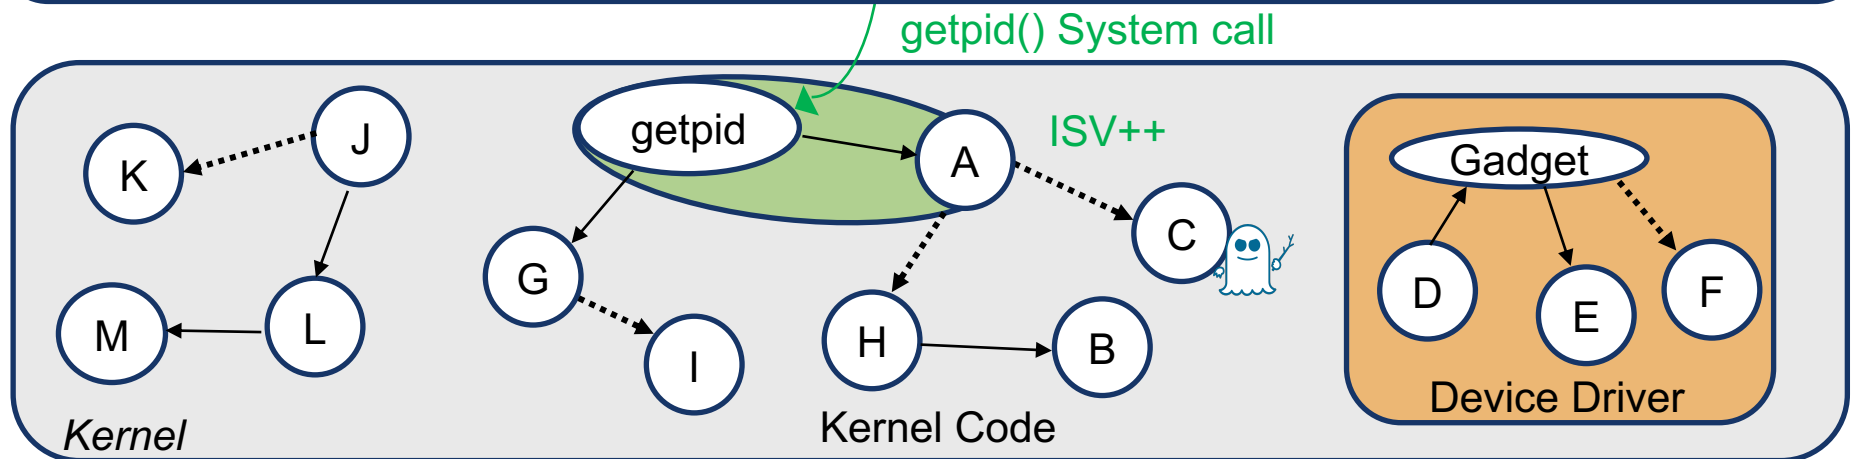
Enhance ISVs: Leverage results of gadget finding to generate stricter ISVs (ISV++)
Swift Mitigation: Quickly exclude vulnerable kernel functions from ISVs



Benefits of ISVs

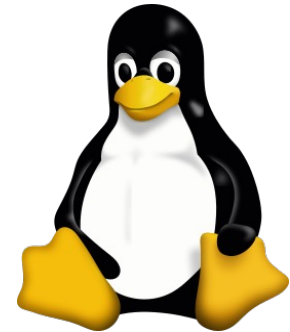


Enhance ISVs: Leverage results of gadget finding to generate stricter ISVs (ISV++)
Swift Mitigation: Quickly exclude vulnerable kernel functions from ISVs

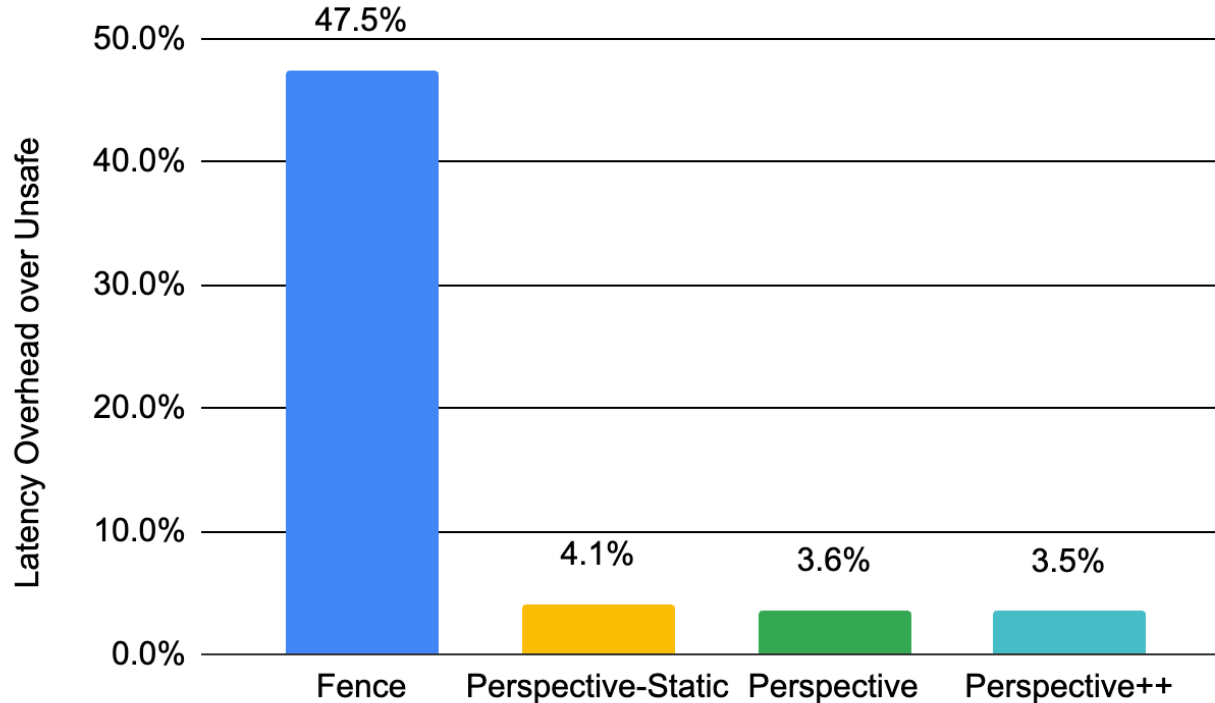


Experiment Setup

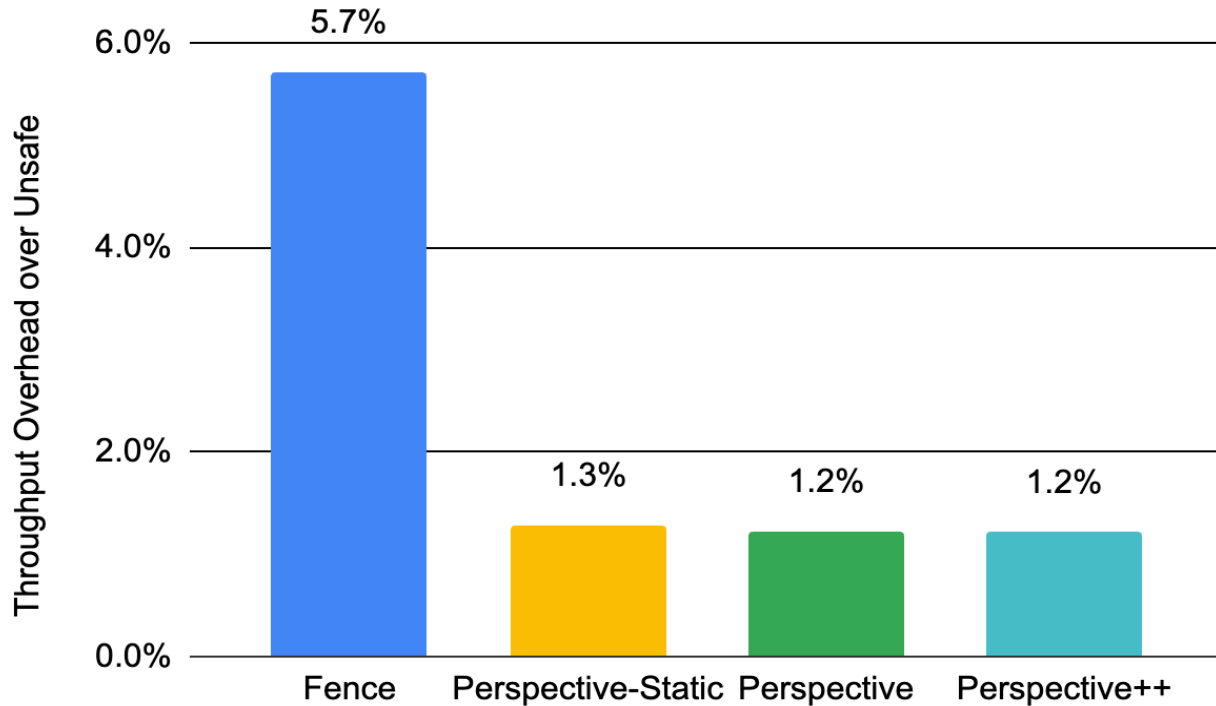
- Gem5 with Linux kernel (v5.4.49) modified for Perspective
- Workloads:
 - LEBench
 - Datacenter Applications -- httpd, nginx, redis, memcached
- Defenses
 - Fence
 - Perspective-Static
 - Perspective
 - Perspective++



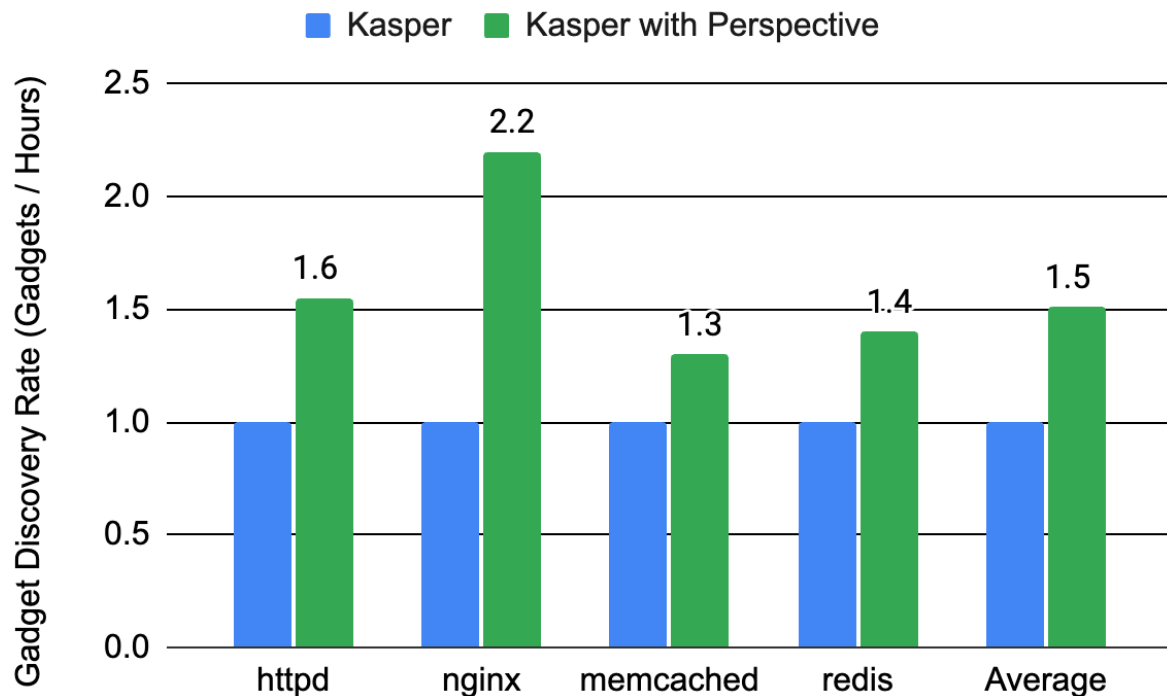
LEBench Evaluation: Average Latency Overhead



Applications Evaluation: Average Throughput Overhead

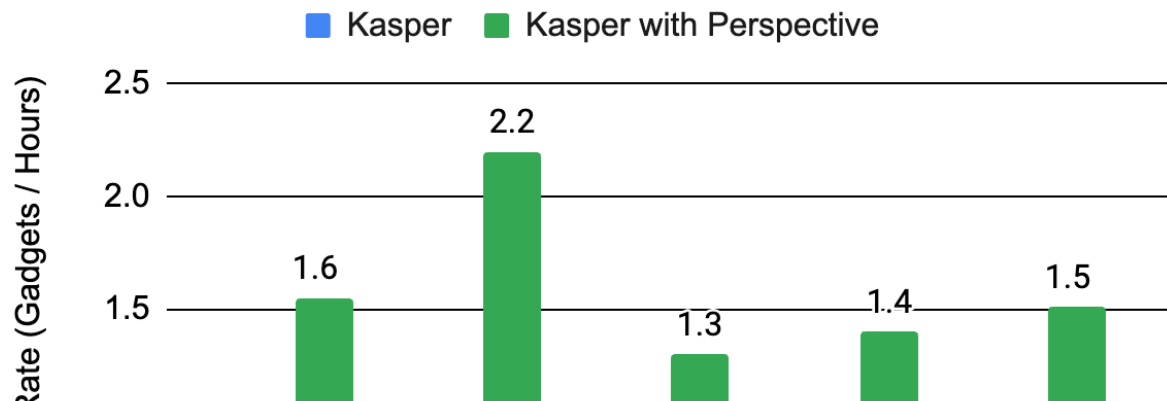


Security Evaluation of ISVs via Kasper¹



¹B. Johannesmeyer, J. Koschel, K. Razavi, H. Bos, and C. Giuffrida, "Kasper: Scanning for Generalized Transient Execution Gadgets in the Linux Kernel," in NDSS, Apr. 2022.

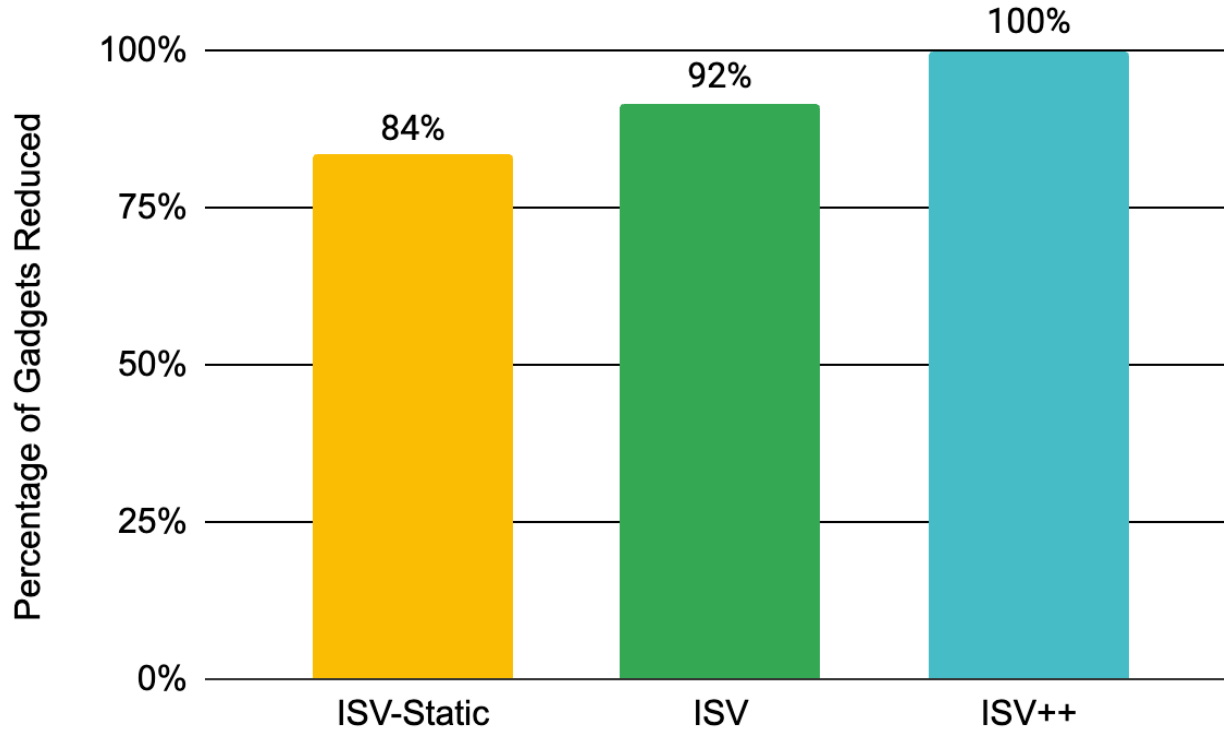
Security Evaluation of ISVs via Kasper¹



ISVs allow for speedup in search for relevant gadgets

¹B. Johannesmeyer, J. Koschel, K. Razavi, H. Bos, and C. Giuffrida, "Kasper: Scanning for Generalized Transient Execution Gadgets in the Linux Kernel," in NDSS, Apr. 2022.

Security Evaluation: Average Gadget Reduction Percent



But Wait There's More!

- Conducted further security evaluations
- Hardware structure evaluations
- Secure kernel memory slab allocator
- ISV hardware caches and virtual address space

Takeaways

- **A taxonomy of transient execution attacks in the kernel**
 - Active and passive attack scenarios
- **Two Perspective OS-HW interfaces \Rightarrow Tailored defense**
 - Data speculation views (DSVs) and Instruction speculation views (ISVs)
- **Security**
 - DSVs completely eliminate active attacks
 - ISVs can
 - Reduce the passive attack surface
 - Accelerate Spectre gadget discovery in the kernel
 - Swiftly mitigate newly discovered gadgets
- **Simple hardware and low execution overhead**
 - Execution overhead of 3.5% on LEBench and 1.2% on datacenter applications