



UNIVERSITÀ DEGLI STUDI DI NAPOLI
FEDERICO II

itee^{PhD}
information technology
electrical engineering



Carmine Cesarano
**Reducing Software Attack Surfaces via
Security Hardening, Fuzzing
and Runtime Enforcement**

Tutor: Roberto Natella

Cycle: XXXVIII

Year: Third

Candidate's information

- MSc degree in **Computer Engineering** (June 2022)
- **Research group:** Dependable and Secure Software Engineering and Real-Time Systems (DESSERT)
- **PhD start-end date:** 01/11/25 – 31/10/2025
- **Scholarship type:** UNINA
- **Periods abroad:** 6 months at KTH Royal Institute of Technology in Stockholm

Summary of study activities

Ad hoc PhD courses

- IoT Data Analysis
- Virtualization Technologies and their Applications
- Statistical Data Analysis for Science and Engineering Research
- Percorso per il rafforzamento delle competenze sulla progettazione europea
- Strategic Orientation for STEM Research & Writing
- Using Deep Learning Properly
- Innovation and Entrepreneurship
- AI Code Generation: Foundations, Evaluation, and Security

Summary of study activities

Conferences attended

- 24th IEEE International Symposium on Software Reliability Engineering (ISSRE 2024)
- 19th European Dependable Computing Conference (EDCC 2024)
- ACM Conference on Computer and Communications Security (CCS2024)
- 36th IEEE International Symposium on Software Reliability Engineering (ISSRE 2025)

Research areas

- **Security Testing** of Industrial Software Systems
 - General-purpose, open-source hypervisors for cloud environments
 - Specialized, closed-source hypervisors for defence
 - Industrial firewall systems for railways
- **Security Hardening** of Container Orchestration Platforms
 - Container Orchestrators like Kubernetes used in cloud
- **Runtime Security** for Open-Source Software Supply Chain
 - Golang programming language widely used for cloud infrastructures

Research results

- Comprehensive analysis of security challenges and mitigation opportunities in edge computing.
- Methodological framework for testing Application Layer Gateways used in critical domains.
- Methodological framework for assessing and certifying isolation properties in partitioning hypervisors.
- Fuzzing framework for testing encrypted network protocols like OpenVPN (**FuEL**).

Thesis

- Fuzzing framework for testing open-source, general purpose hypervisors (**IRIS**).
- Fuzzing framework for testing closed-source, industrial hypervisors (**FuzzBox**).
- API filtering system for reducing the attack surface of Kubernetes (**KubeFence**)
- Language-specific taxonomy of supply chain attack vectors in Go and static analysis tool for identifying them (**GoSurf**)
- Runtime enforcement system mitigating supply chain attacks in Go (**GoLeash**)

Research products (1)

[P1]	<p><u>C. Cesarano</u>, D. Cotroneo, L. De Simone Towards Assessing Isolation Properties in Partitioning Hypervisors 33rd IEEE International Symposium on Software Reliability Engineering (ISSRE2022) Charlotte, NC, USA, 31 October–03 November 2022, IEEE, DOI: 10.1109/ISSREW55968.2022.00067</p>
[P2]	<p><u>C. Cesarano</u>, M. Cinque, D. Cotroneo, L. De Simone, G. Farina IRIS: a Record and Replay Framework to Enable Hardware-assisted Virtualization Fuzzing 53rd IEEE/IFIP International Conference on Dependable Systems and Networks (DSN2023) Porto, Portugal, 27-30 June 2023, IEEE, DOI: 10.1109/DSN58367.2023.00045</p>
[P3]	<p><u>C. Cesarano</u> Security Assessment and Hardening of Fog Computing Systems 34th IEEE International Symposium on Software Reliability Engineering (ISSRE2023) Florence, Italy, 09-12 October 2023, IEEE, DOI: 10.1109/ISSREW60843.2023.00037</p>
[P4]	<p><u>C. Carmine</u>, R. Natella Securing an Application Layer Gateway: An Industrial Case Study 19th European Dependable Computing Conference (EDCC2024) Leuven, Belgium, 08-11 April 2024, IEEE, DOI: 10.1109/EDCC61798.2024.00025</p>
[P5]	<p><u>C. Cesarano</u>, V. Andersson, R. Natella, M. Monperrus GoSurf: Identifying Software Supply Chain Attack Vectors in Go ACM Workshop on Software Supply Chain Offensive Research and Ecosystem Defenese (SCORED24) Salt Lake City, Utah, USA, 14-18 October 2024, ACM, DOI: 10.1145/3689944.3696166</p>

Research products (2)

[P6]	<p>C. Cesarano, A. Foggia, G. Roscigno, L. Andreani, R. Natella GENIO: Synergizing Edge Computing with Optical Network Infrastructures IEEE Communication Magazine (COMMAG-IEEE) Volume 63, Issue 7, Pages: 154-160, July 2025, IEEE, DOI: 10.1109/MCOM.002.2400382</p>
[P7]	<p>C. Cesarano, R. Natella KubeFence: Security Hardening of the Kubernetes Attack Surface. 55th Annual IEEE/IFIP International Conference on Dependable Systems and Networks (DSN) Naples, Italy, 23-26 June 2025, IEEE, DOI: 10.1109/DSN64029.2025.00054</p>
[P8]	<p>C. Cesarano, A. Foggia, G. Roscigno, L. Andreani, R. Natella Security-by-Design at the Telco Edge with OSS: Challenges and Lessons Learned 55th Annual IEEE/IFIP International Conference on Dependable Systems and Networks (DSN) Naples, Italy, 23-26 June 2025, IEEE, DOI: 10.1109/DSN64029.2025.00054</p>
[P9]	<p>C. Cesarano, R. Natella FuzzBox: Blending Fuzzing into Emulation for Binary-Only Embedded Targets Springer Cybersecurity Journal (accepted, to appear)</p>

PhD thesis overview

Problem statement

Modern software systems expose large, multi-layered attack surfaces, that current methods fail to effectively reduce

Objective

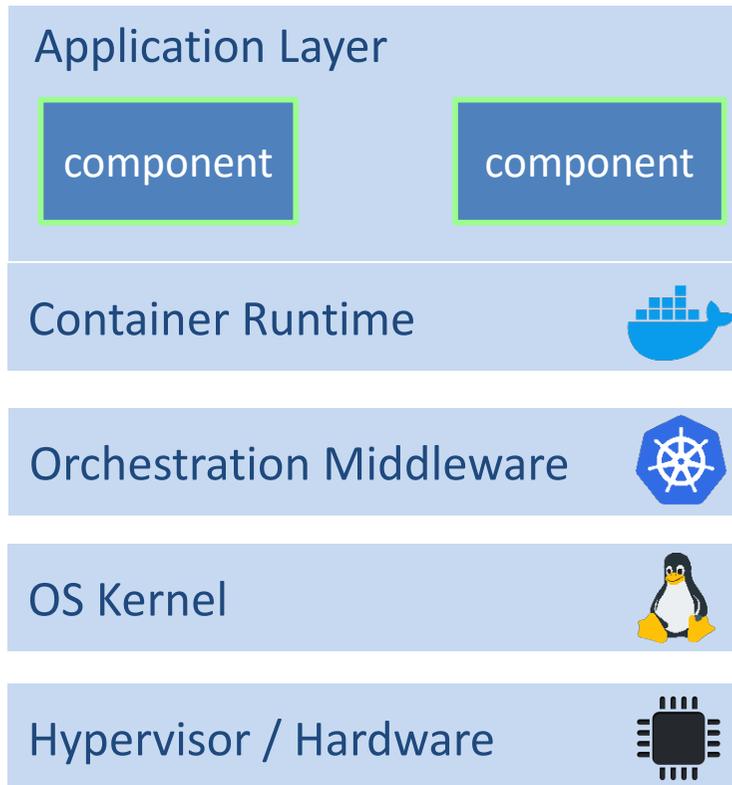
Develop cross-layer, fine-grained, and automated techniques for attack surface reduction that enhance system security without compromising functionality

Methodology

- Characterize attack surfaces and identify layer-specific weaknesses
- Design and prototype five complementary techniques combining static and dynamic analysis
- Evaluate their effectiveness through real-world deployments, benchmarks, vulnerability discovery, and attack mitigation

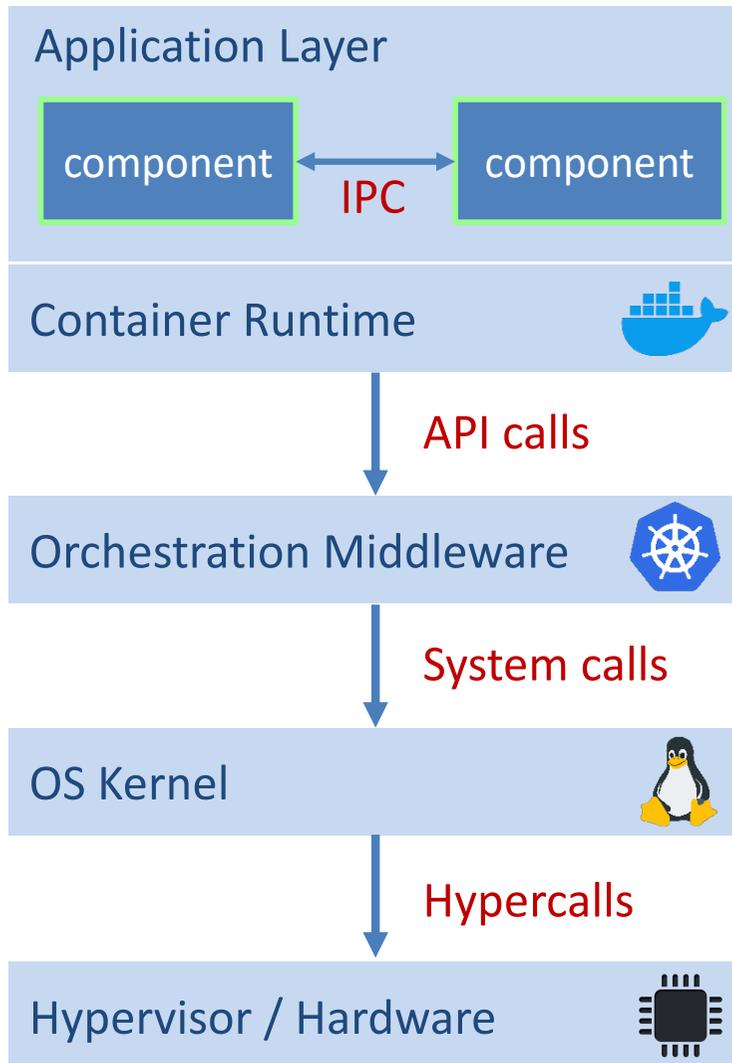
Context

Modern Software Stacks: Complexity and Exposure



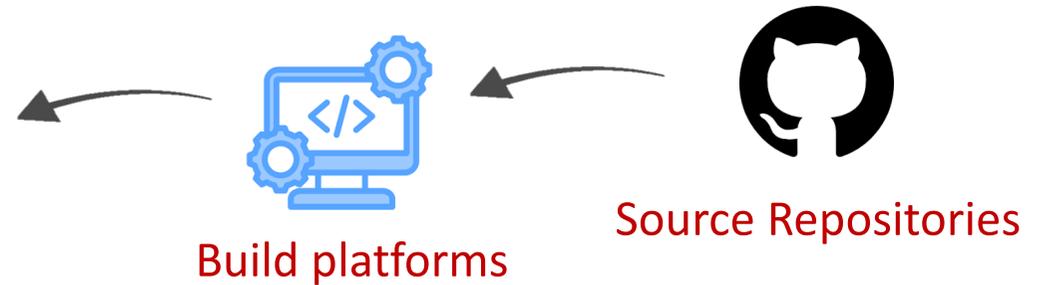
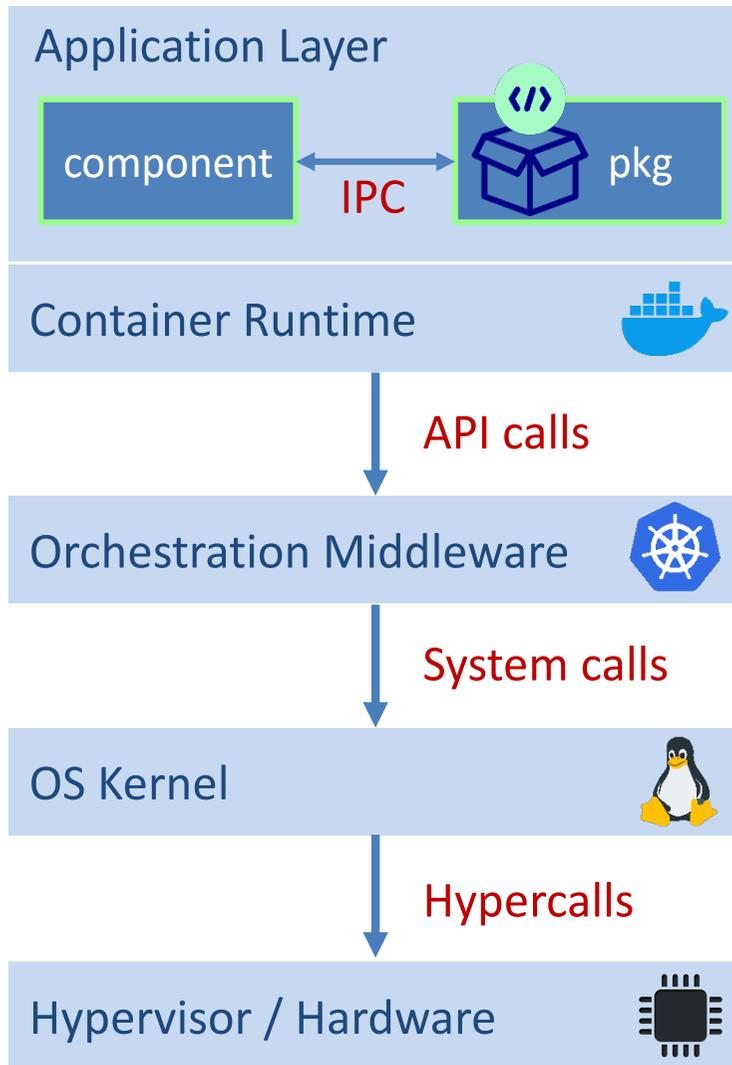
Context

Modern Software Stacks: Complexity and Exposure



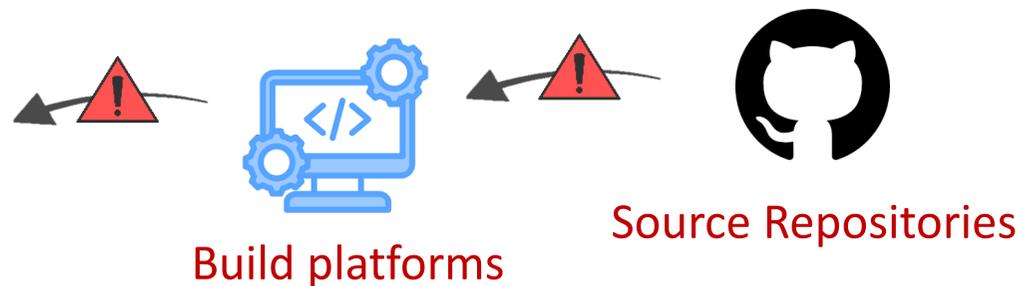
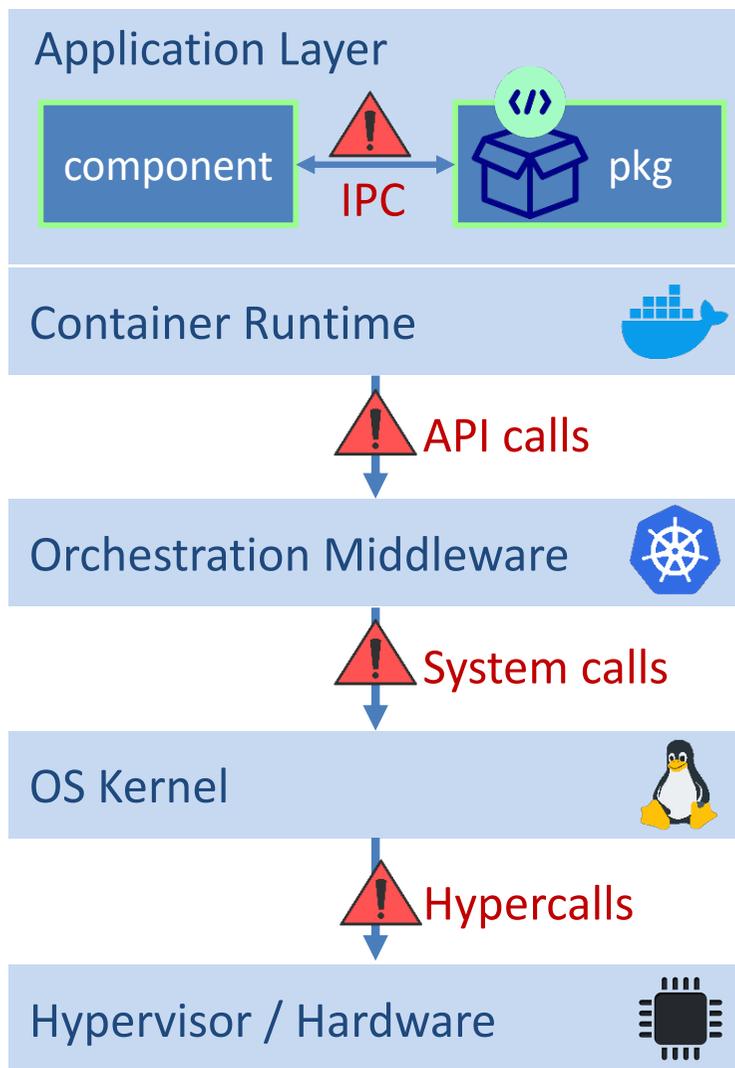
Context

Modern Software Stacks: Complexity and Exposure



Research Problem

Modern Software Stacks: Complexity and Exposure



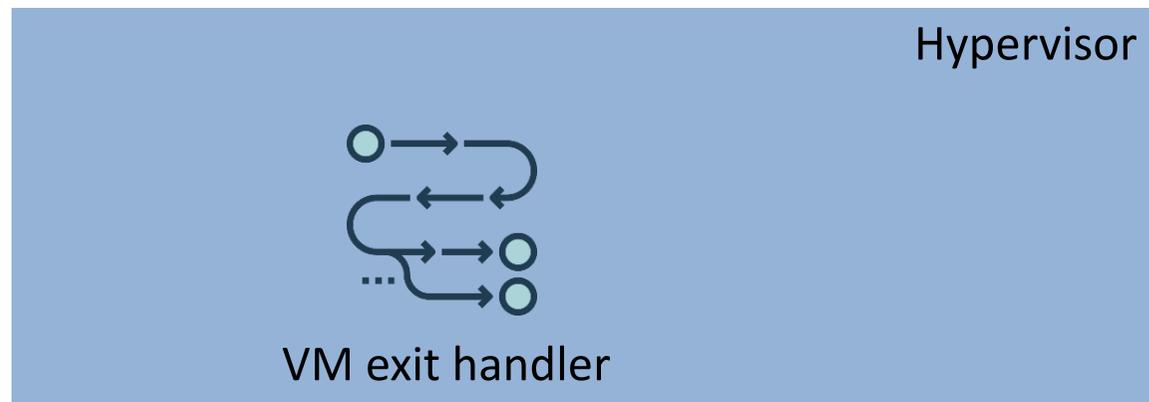
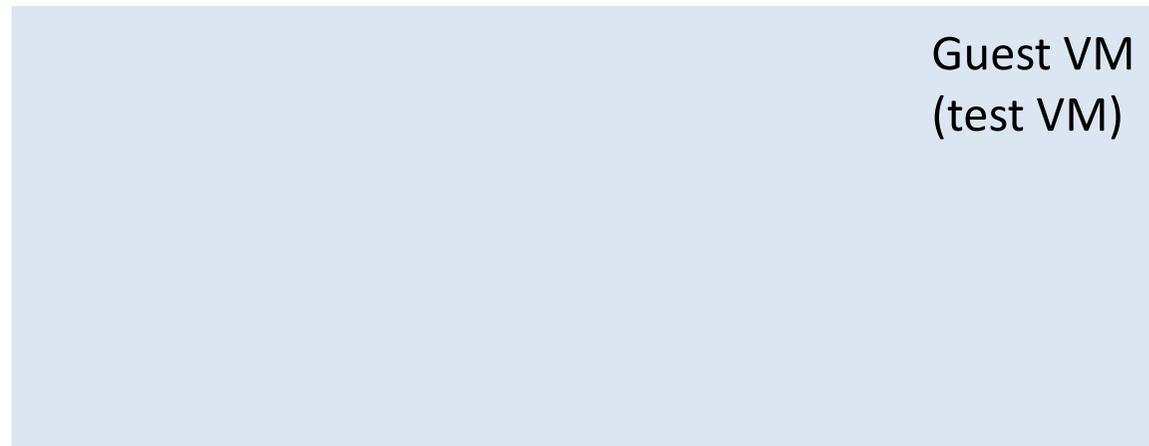
Interfaces represent attack surfaces

How can we systematically reduce them?

Contribution 1

IRIS: Fuzzing the Hypervisor Surface

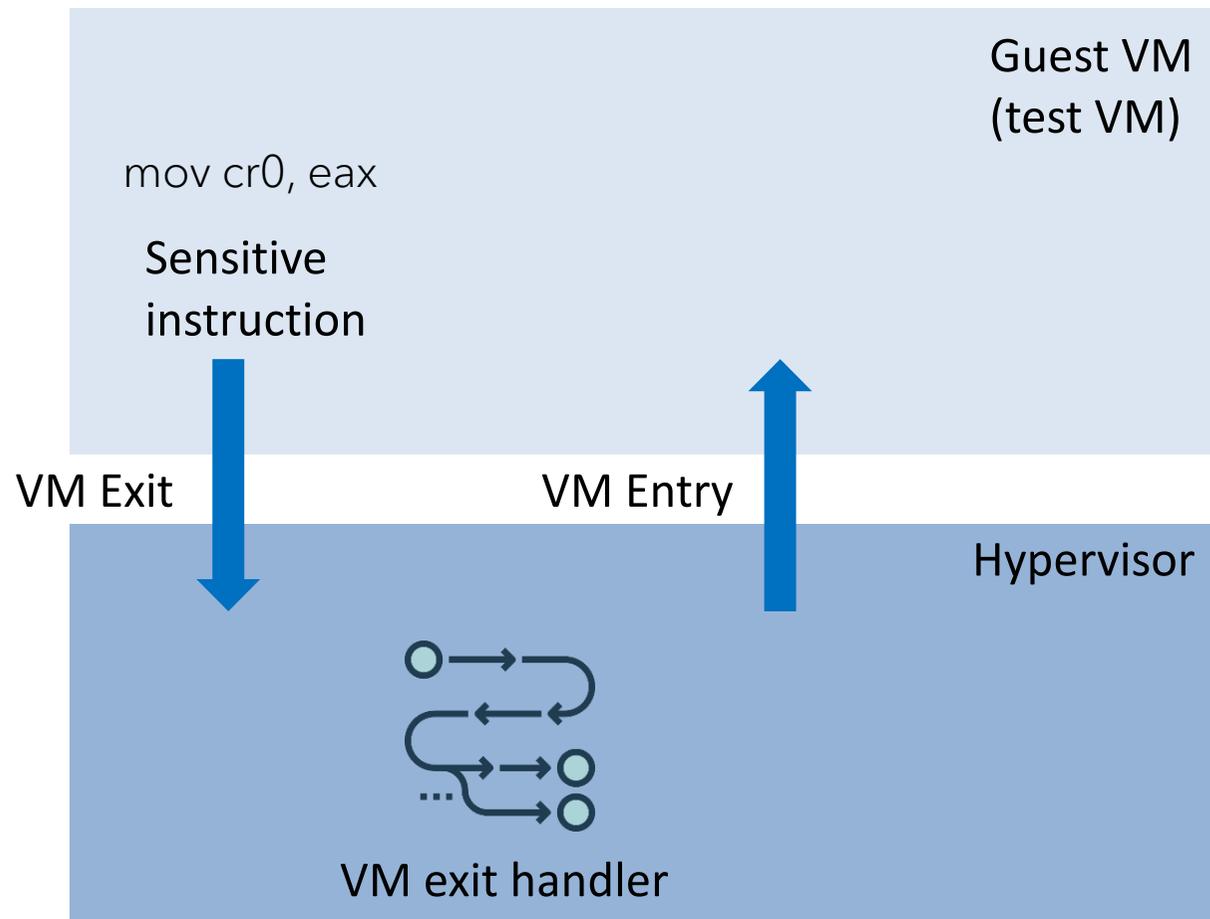
Testing hypervisor logic is stateful and hard



Contribution 1

IRIS: Fuzzing the Hypervisor Surface

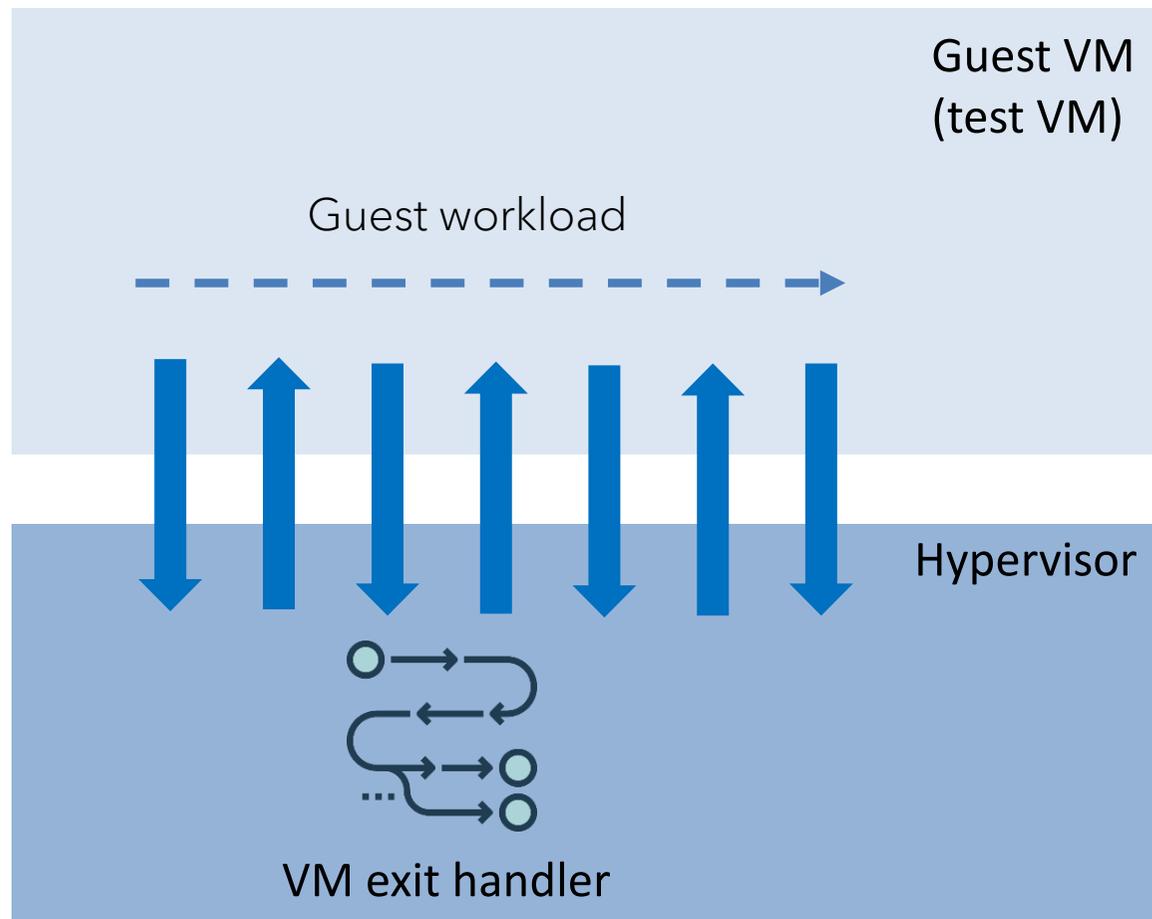
A single VM exit handling depends on the guest VM state, sensitive instruction invoked, hypervisor current state and previous VM exits



Contribution 1

IRIS: Fuzzing the Hypervisor Surface

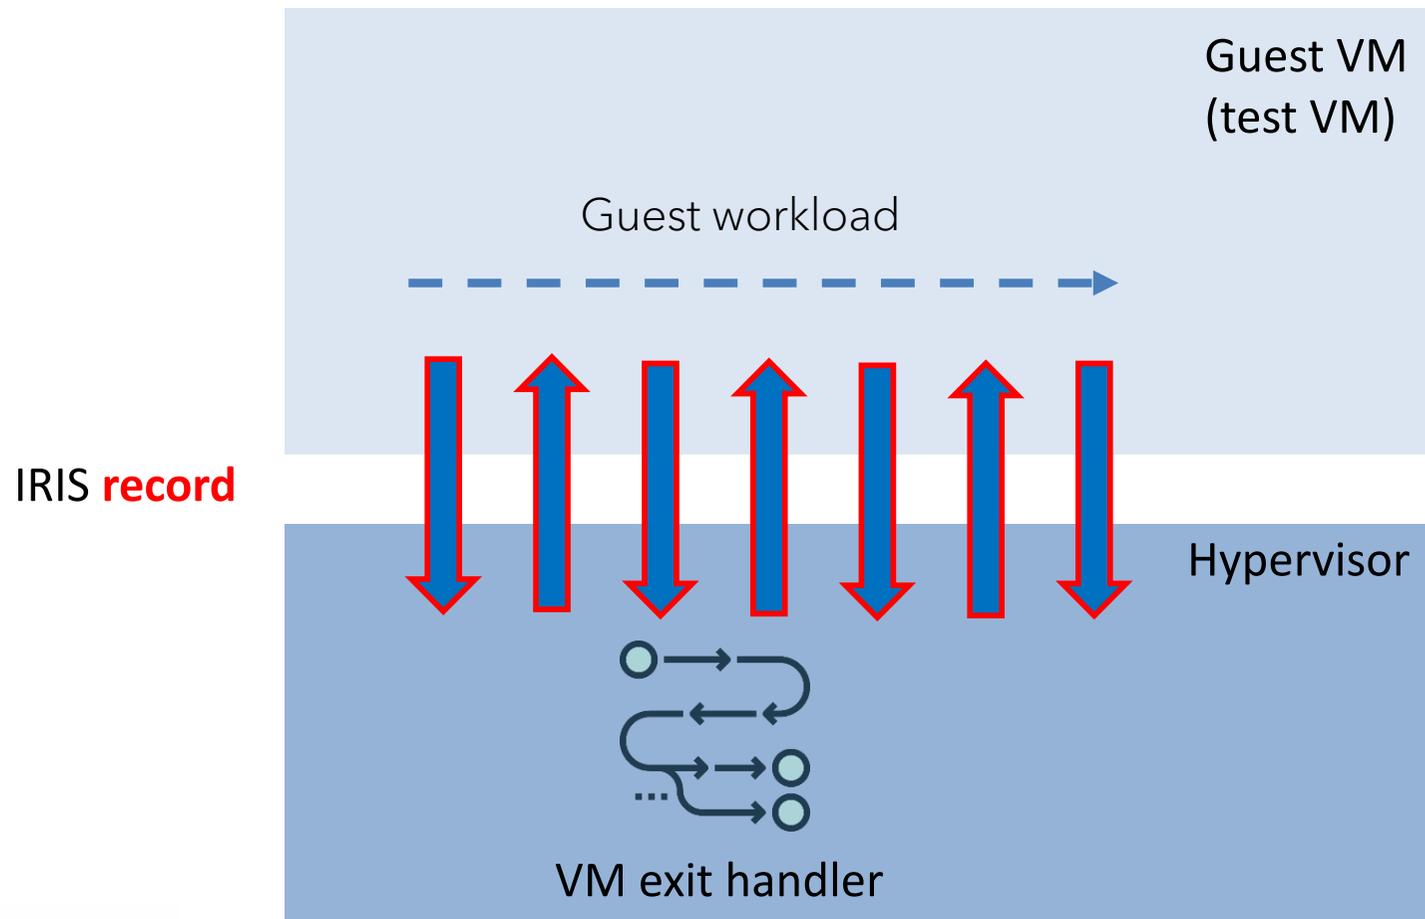
Reproducing that sequence from a test guest requires re-implementing OS-level flows (prohibitive)



Contribution 1

IRIS: Fuzzing the Hypervisor Surface

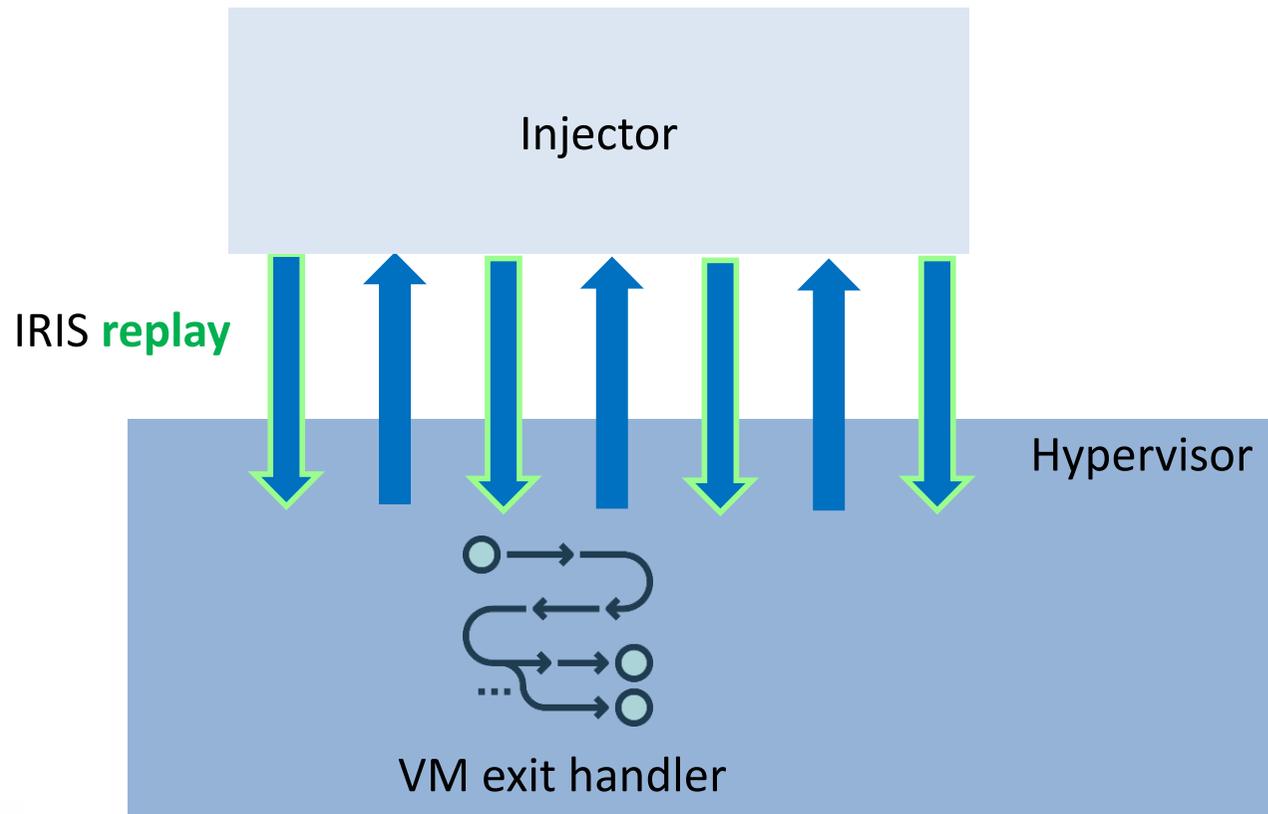
IRIS record VM exit events during hypervisor intervention by executing a nominal guest workload



Contribution 1

IRIS: Fuzzing the Hypervisor Surface

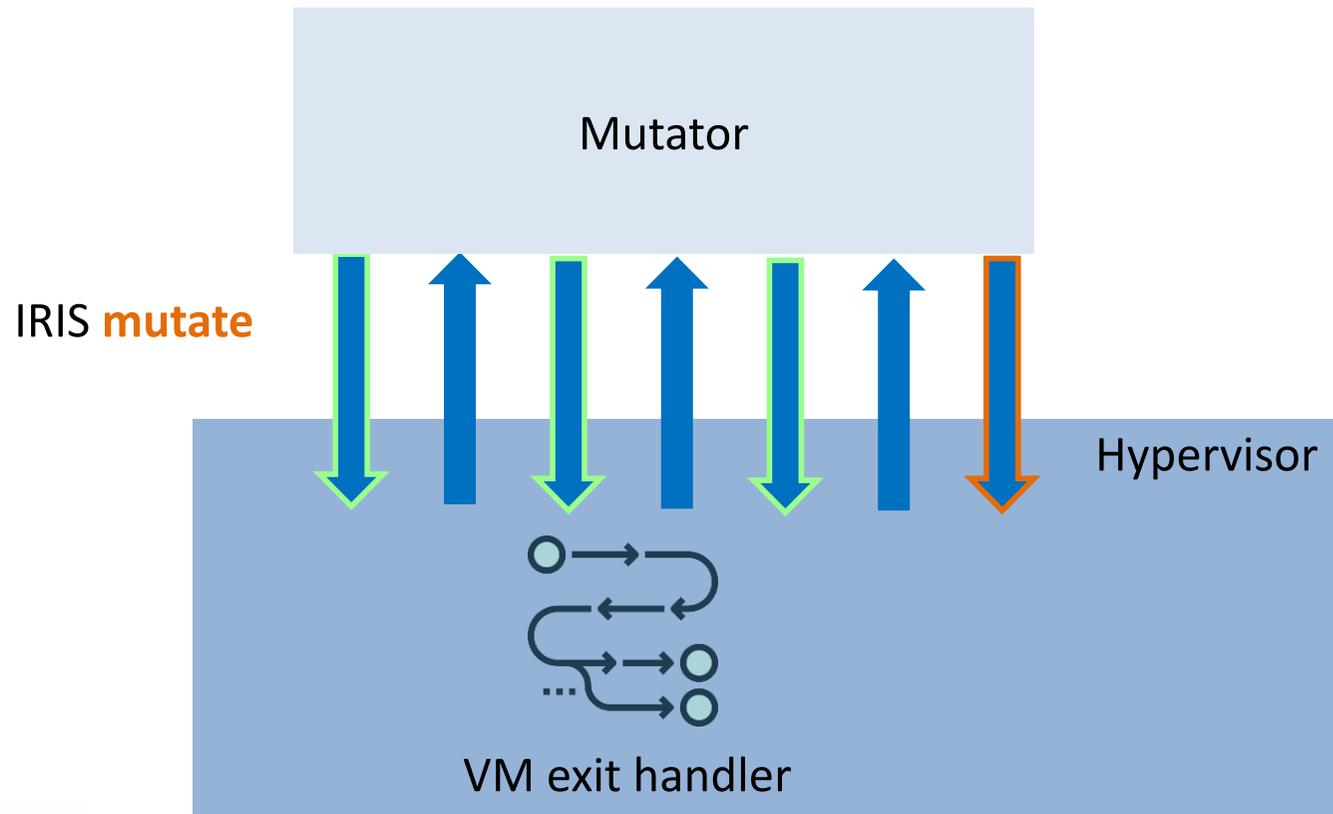
IRIS replay such events to quickly reach deep hypervisor states



Contribution 1

IRIS: Fuzzing the Hypervisor Surface

IRIS can mutate VM exits in the replayed sequence to test corner cases, facilitating fuzzing of the hypervisor logic



Contribution 1

IRIS: Fuzzing the Hypervisor Surface

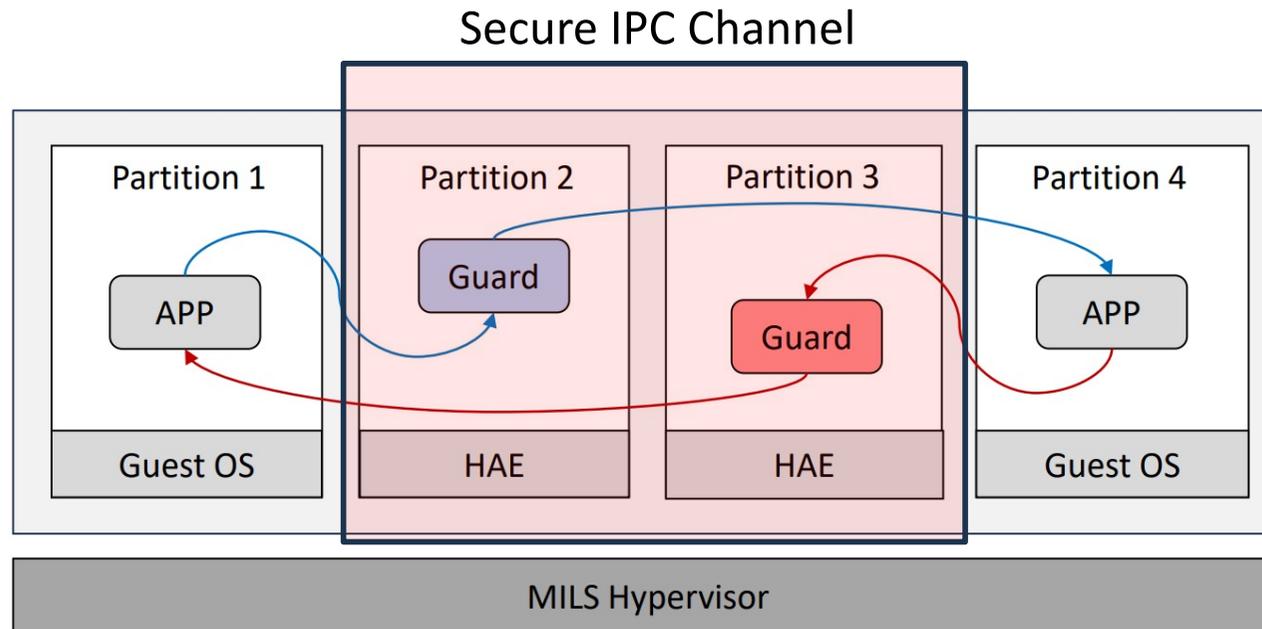
Evaluations on three real guest workloads:

- Reproduces valid hypervisor states identical to real guest executions by replaying VM exit sequences
- Matches 92.2-100% of code coverage of real guest workload
- Achieves 42.5–99.6% faster execution time compared to full guest runs

Contribution 2

FuzzBox: Fuzzing the Syscall/IPC Surface

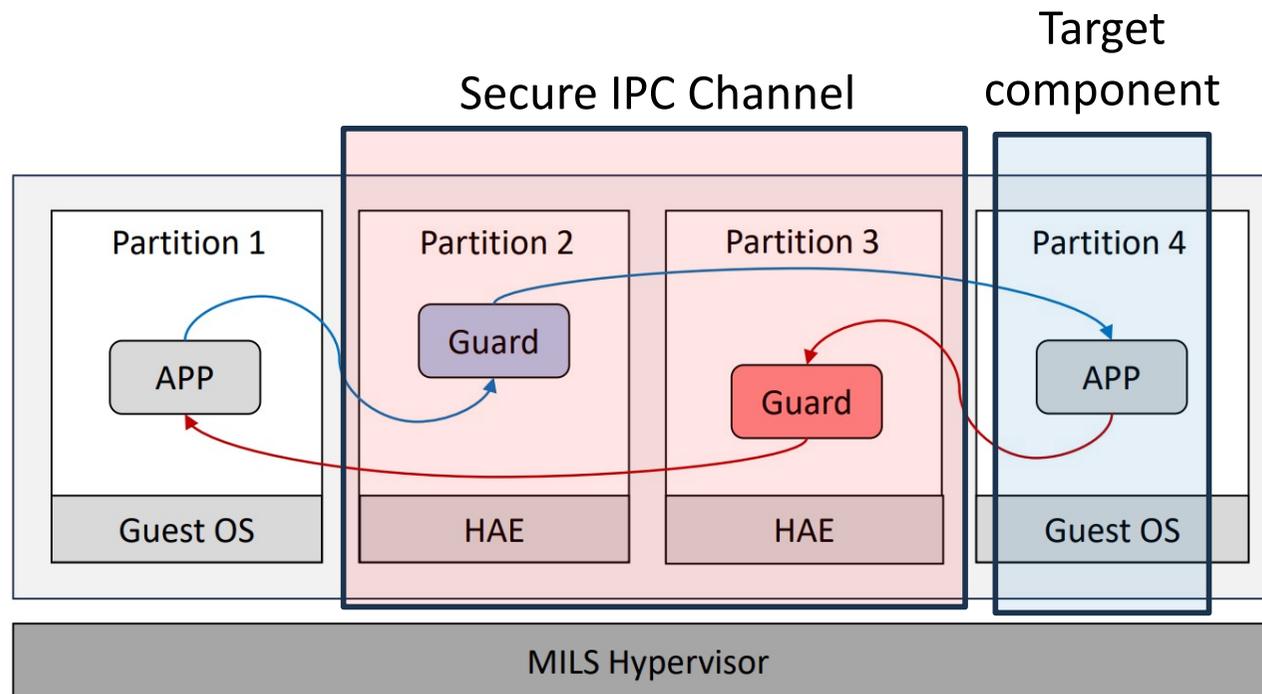
Testing internal opaque components within closed-source binaries is challenging



Contribution 2

FuzzBox: Fuzzing the Syscall/IPC Surface

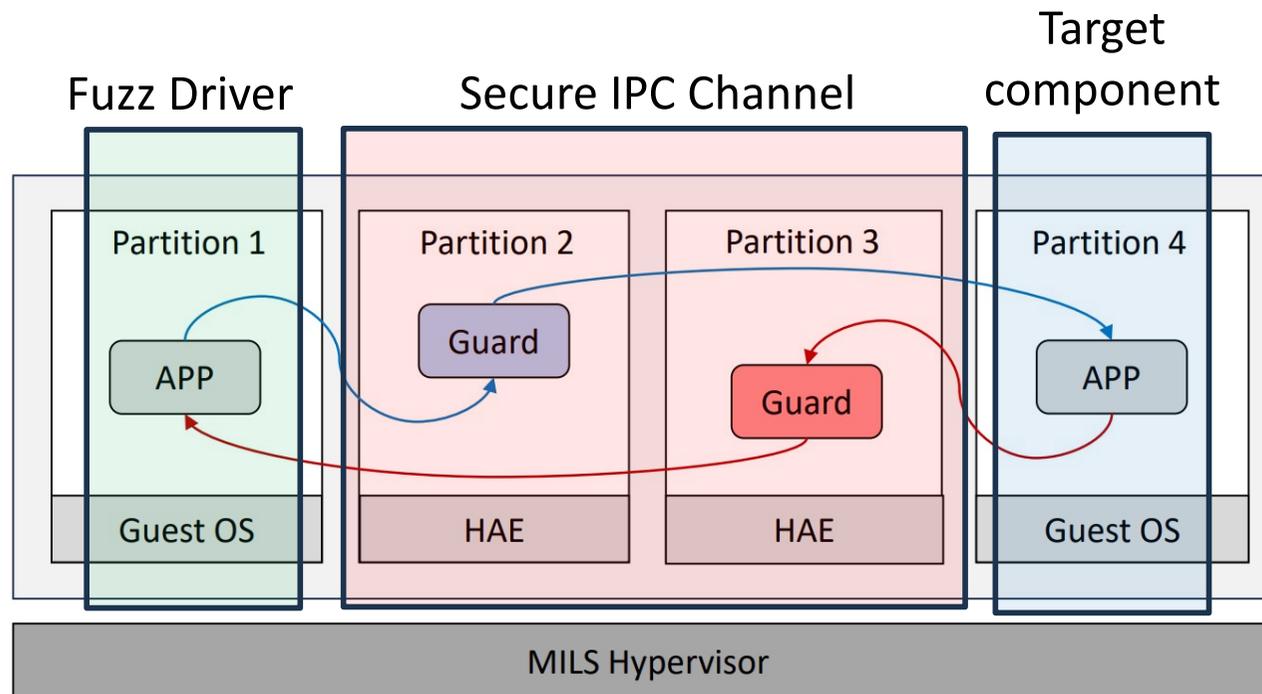
Testing internal opaque components within closed-source binaries is challenging



Contribution 2

FuzzBox: Fuzzing the Syscall/IPC Surface

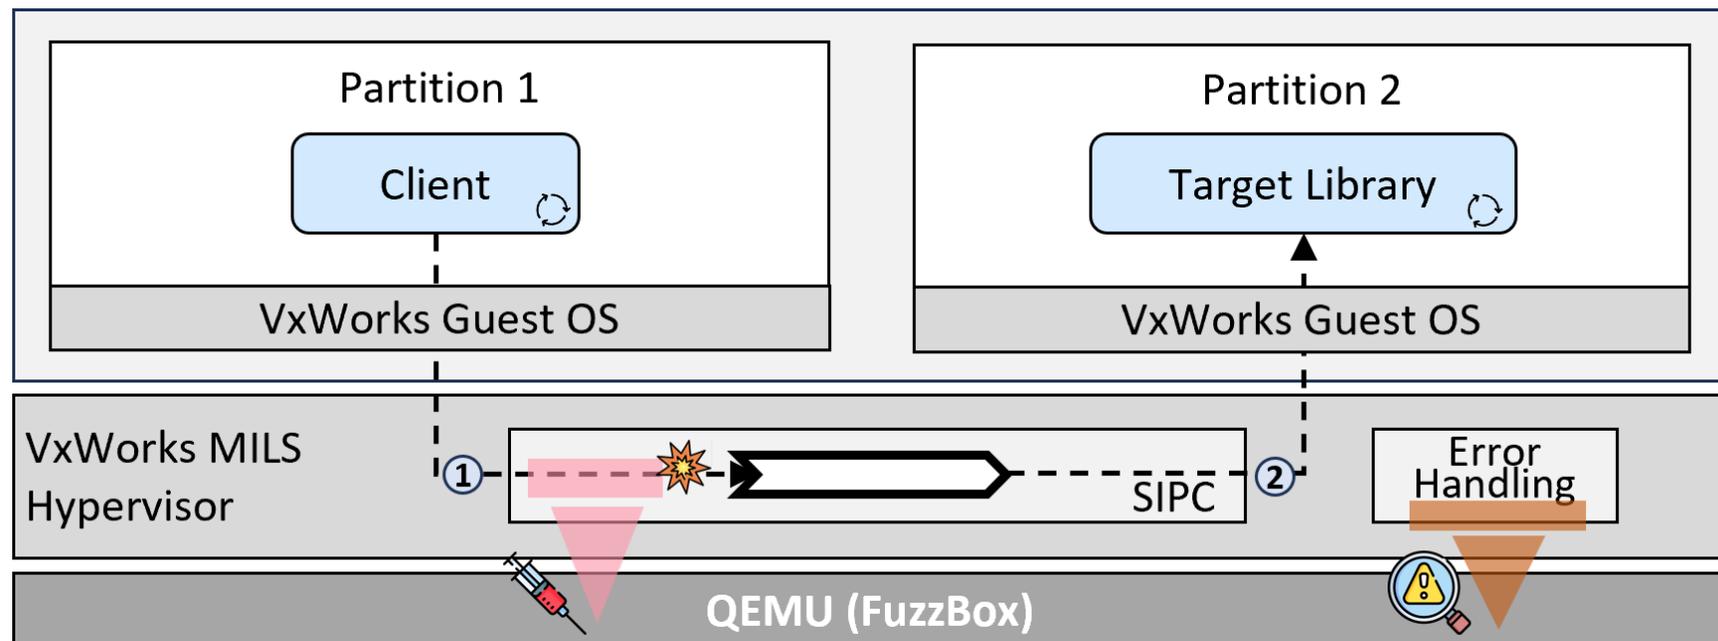
Challenges of Intrusiveness, Toolchain Dependency and Hardware Dependency make fuzzing of industrial, closed-source binaries infeasible



Contribution 2

FuzzBox: Fuzzing the Syscall/IPC Surface

FuzzBox integrate fuzzing with emulation to intercept and mutate internal function calls, gather branch coverage transparently, and monitor error handling



Contribution 2

FuzzBox: Fuzzing the Syscall/IPC Surface

Before FuzzBox fuzzing internal components in MILS-based systems required:

- Adding a fuzz driver inside a partition
- Rebuilding the entire system binary
- Performing black-box fuzzing

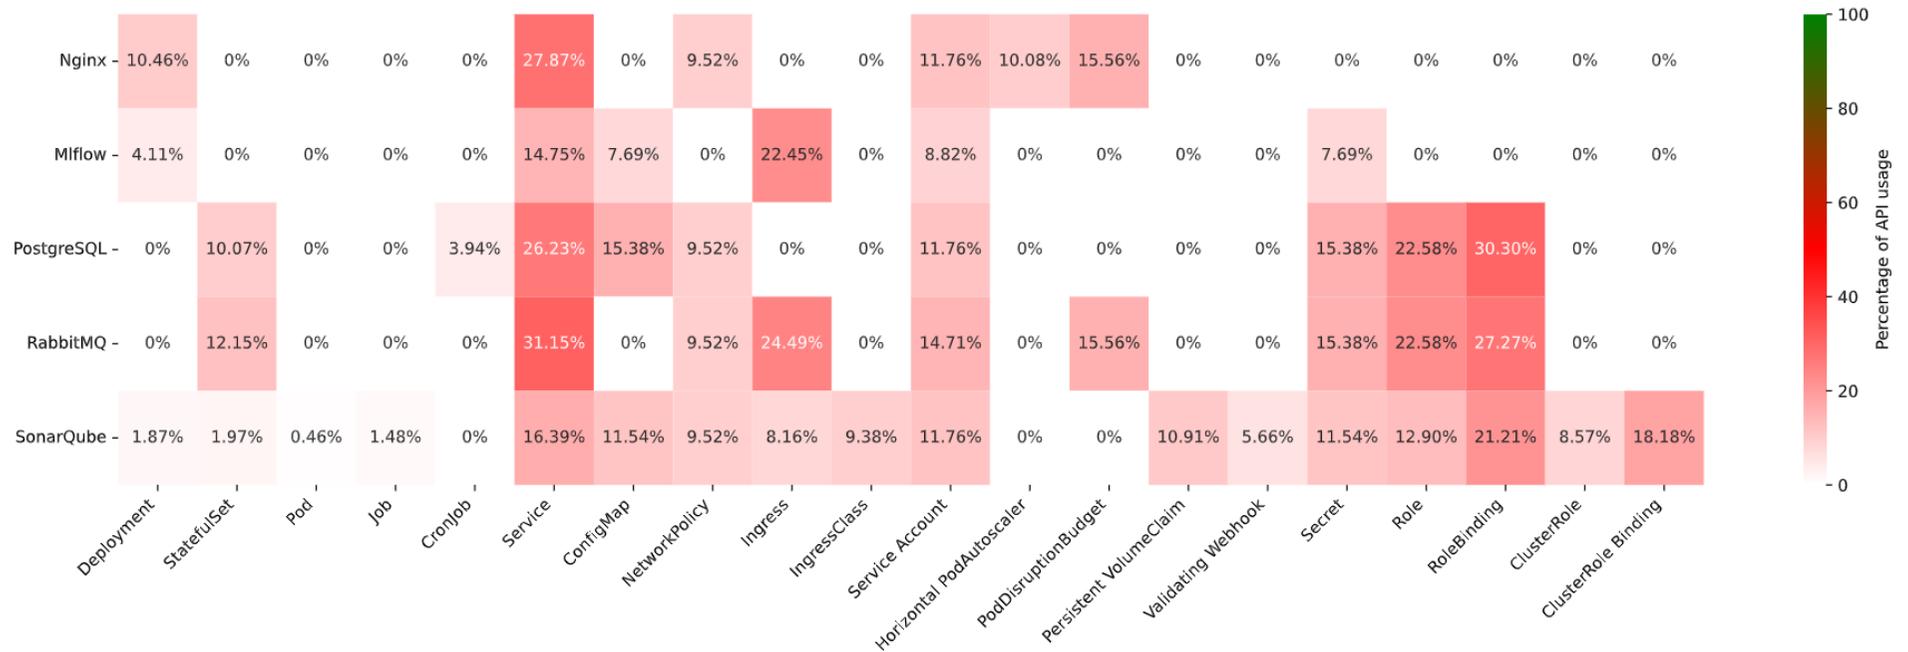
With FuzzBox

- Enabled direct gray-box fuzzing of internal components without rebuilds
- 2X higher fuzzing throughput
- ~58% coverage improvement
- Faster bug discovery

Contribution 3

KubeFence: Hardening the Orchestration Surface

Most of the Kubernetes API is unused by still attackable



Contribution 3

KubeFence: Hardening the Orchestration Surface

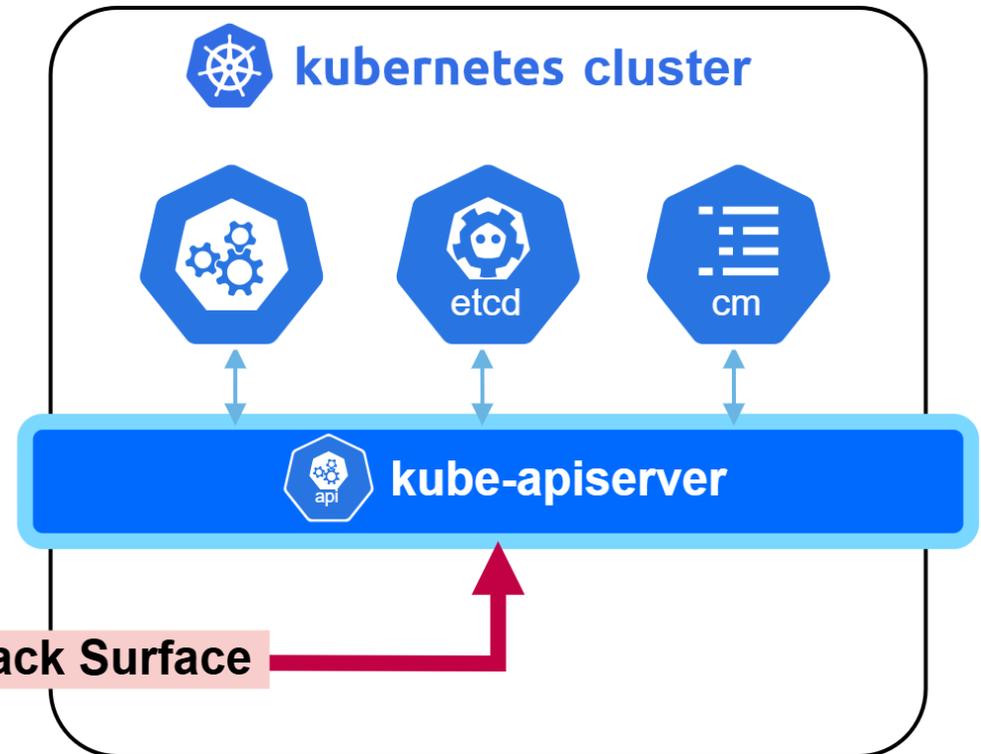
Current methods to harden orchestration APIs (e.g., RBAC) are coarse-grained

API Request

```
POST /api/v1/namespaces/default/pod
Host: <api-server-url>
Content-Type: application/json
Accept: application/json
{
  "apiVersion": "v1",
  "kind": "Pod",
  "metadata": {
    "name": "example-pod",
  },
  "spec": {
    "hostNetwork": true,
    "containers": [...]
```



Large Attack Surface

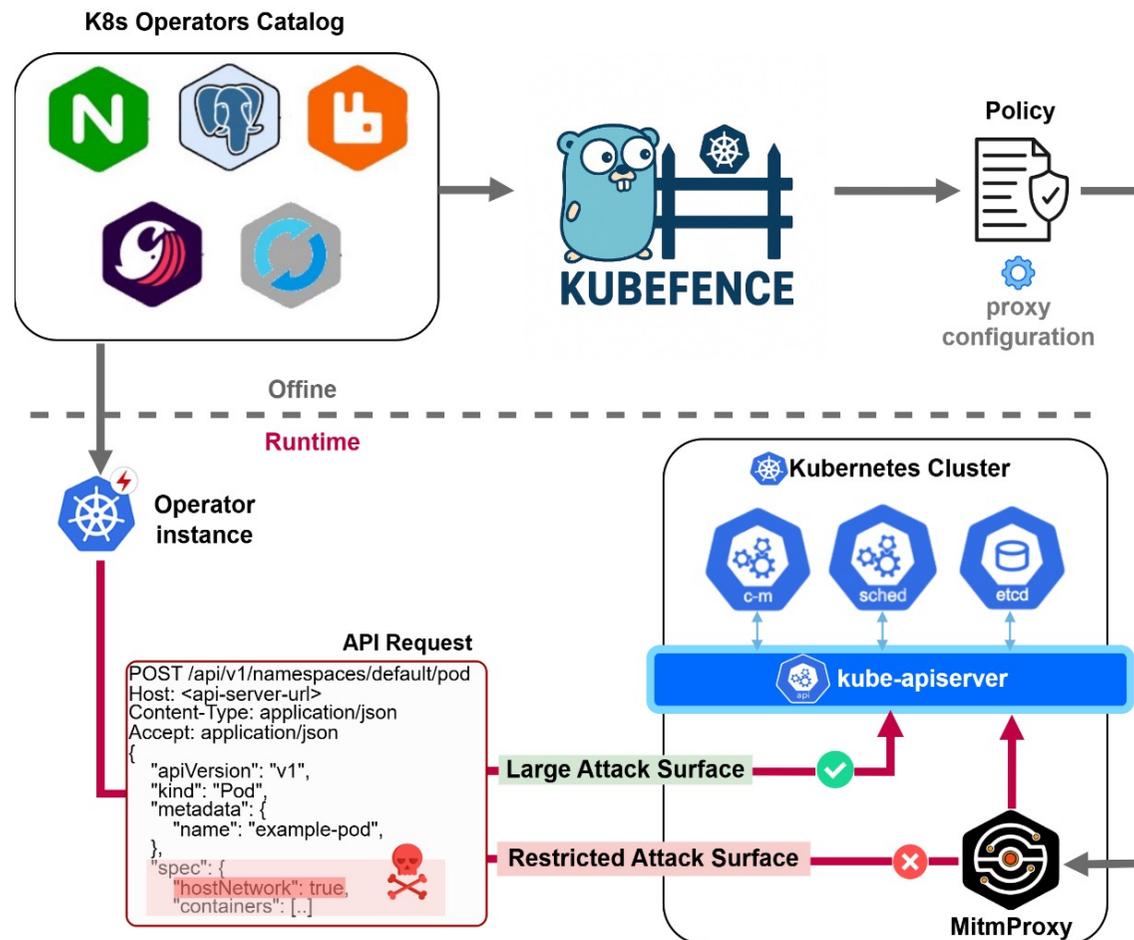


RBAC stops resource kind **not specs!**

Contribution 3

KubeFence: Hardening the Orchestration Surface

KubeFence learns access policies at the level of spec, for a specific K8s app, and apply runtime enforcement allowing only required APIs fields and values



Contribution 3

KubeFence: Hardening the Orchestration Surface

KubeFence vs. RBAC

- KubeFence achieved an average 35% reduction in cluster attack surface compared to RBAC
- Mitigated 8 CVE exploits and 7 misconfigurations, while RBAC prevented none.
- Introduced only ~21% overhead (~50 ms). Overhead applies only to cluster management operations. No impact on workload performance.

Contribution 4

GoSurf: Identifying Software Supply Chain Attack Vectors

How can attackers hide the execution of malicious code in Go dependencies?

Automatic Behavior

- Static Code Generators
- Testing Functions

Pre-Build Phase

Unexpected Behaviors

- Global Variable Initialization
- Initialization Hooks
- Constructors Methods

Initialization Phase

Dynamic Behaviors

- Reflection
- Interfaces
- Unsafe Pointers
- CGO Functions
- Assembly Functions
- Dynamic Plugins
- External Execution

Execution Phase

We created a taxonomy of 12 attack vectors across three stages of the Go package lifecycle

Contribution 4

GoSurf: Identifying Software Supply Chain Attack Vectors

How can developers prioritize code reviews of imported packages and/or understand if packages added vectors across new versions?

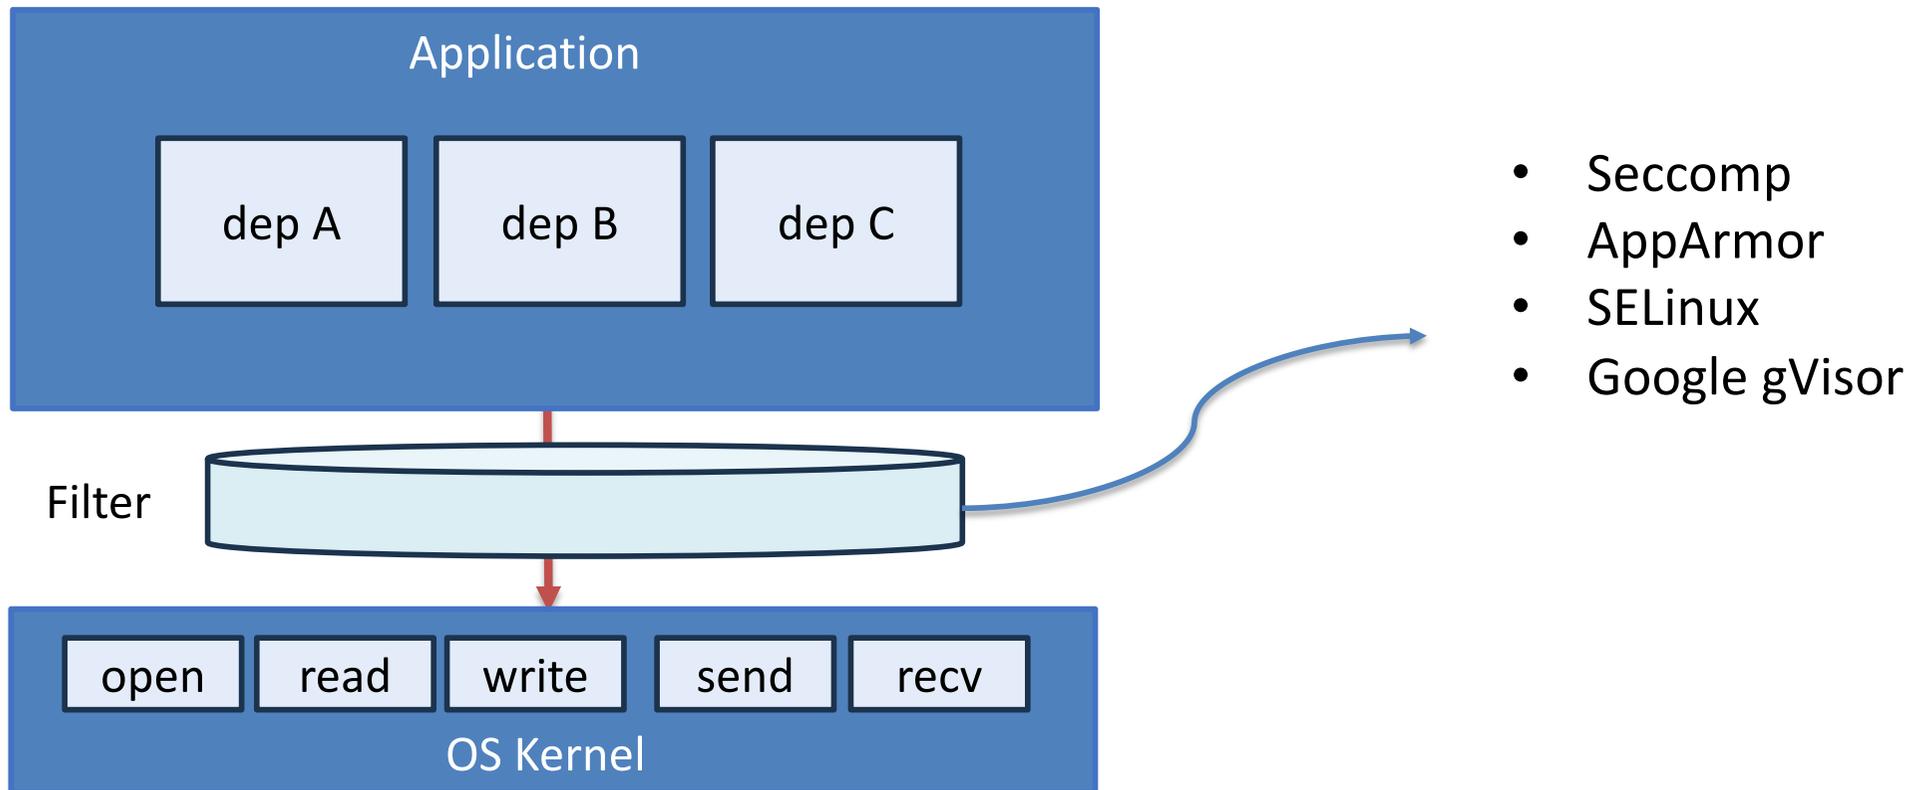
Line	Version	Category	Count
1	v1.14.5	Attack Surface Analysis: geth-v1.14.5	-
2	v1.14.5	[P1] Static Code Generation:	34
3	v1.14.5	[P2] Testing Functions:	2132
4	v1.14.5	[I1] Global Variable Initialization:	1116
5	v1.14.5	[I2] init() Functions:	76
6	v1.14.5	[E1] Constructor Methods:	6470
7	v1.14.5	[E2] Reflection:	128
8	v1.14.5	[E3] Interfaces:	6343
9	v1.14.5	[E4] Unsafe Pointers:	26
10	v1.14.5	[E5] CGO Functions:	21
11	v1.14.5	[E6] Assembly Functions:	12
1	v1.14.6	Attack Surface Analysis: geth-v1.14.6	+
2	v1.14.6	[P1] Static Code Generation:	35
3	v1.14.6	[P2] Testing Functions:	2134
4	v1.14.6	[I1] Global Variable Initialization:	1133
5	v1.14.6	[I2] init() Functions:	72
6	v1.14.6	[E1] Constructor Methods:	6468
7	v1.14.6	[E2] Reflection:	132
8	v1.14.6	[E3] Interfaces:	6377
9	v1.14.6	[E4] Unsafe Pointers:	21
10	v1.14.6	[E5] CGO Functions:	21
11	v1.14.6	[E6] Assembly Functions:	12

GoSurf makes actionable the taxonomy finding occurrences of the 12 attack vectors in the code through AST analysis

Contribution 5

GoLeash: Mitigating Software Supply Chain Attacks at Runtime

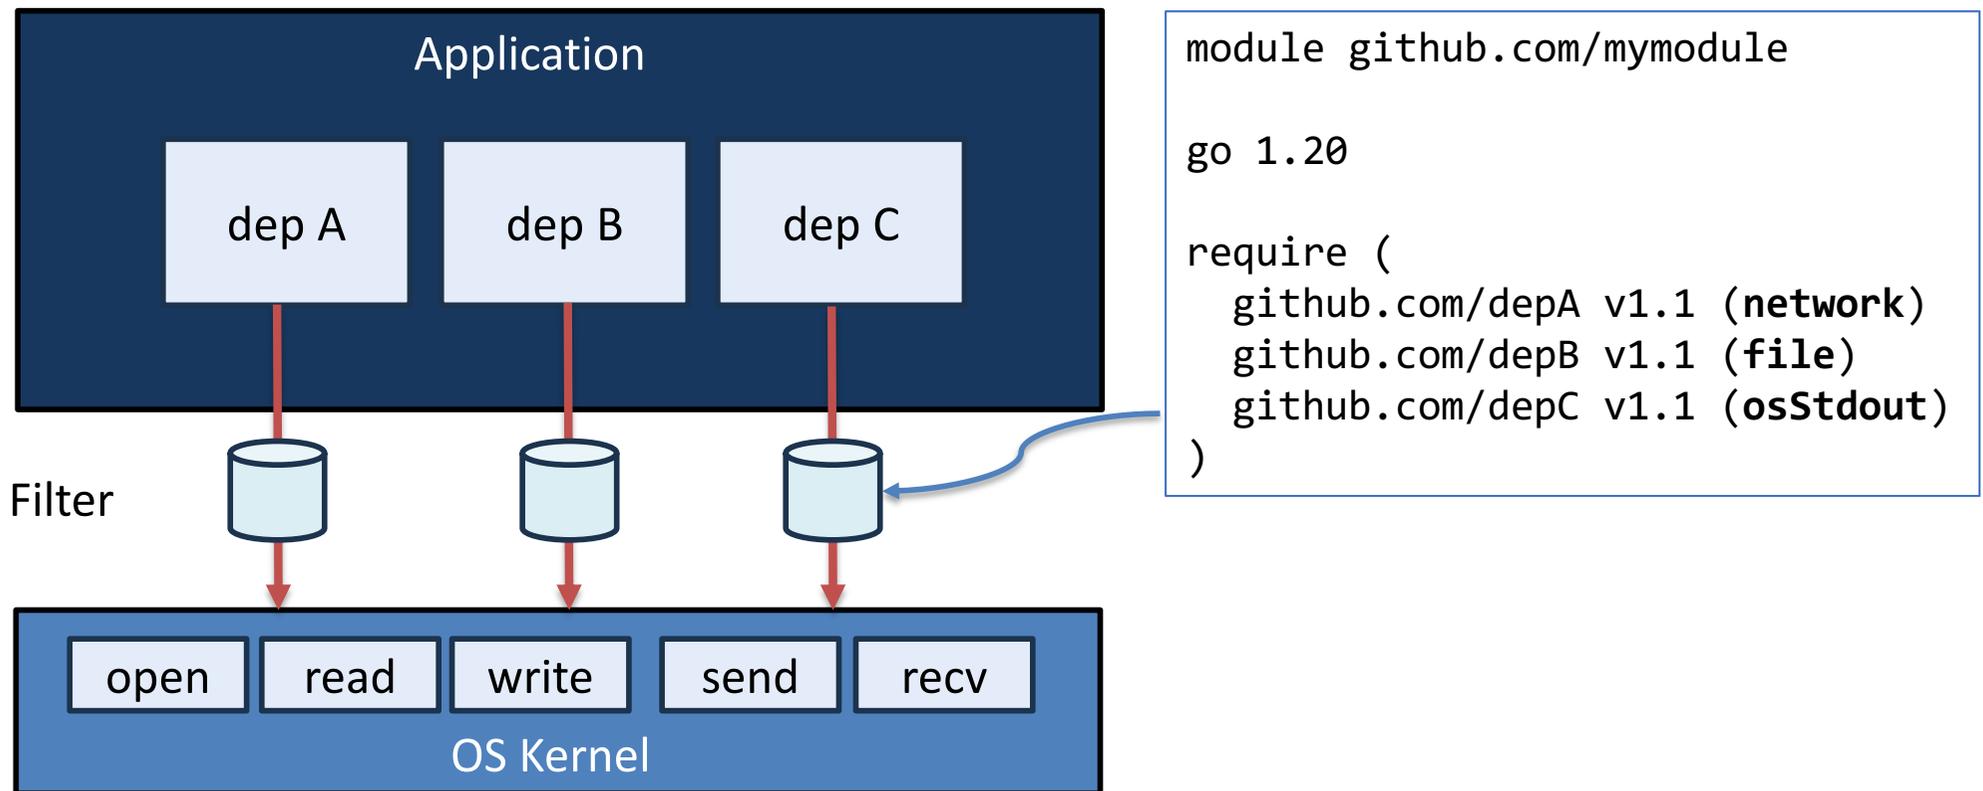
Application-level sandboxing fails to restrict capabilities of untrusted open source packages that are built within our application



Contribution 5

GoLeash: Mitigating Software Supply Chain Attacks at Runtime

GoLeash applies package-level enforcement for allowed capabilities at runtime, mitigating supply chain attacks



Contribution 5

GoLeash: Mitigating Software Supply Chain Attacks at Runtime

Evaluated on 5 real-world Go projects (Kubernetes, etcd, CoreDNS, frp, geth)

- Detected 98% of 3,365 injected supply chain attacks vs 32% with baseline process-level enforcement
- Robust to obfuscation, maintained 92-100% detection across 5,224 obfuscated attacks
- Outperformed Capslock (static analysis) by up to +70 pp detection rate
- Average runtime overhead: **~9%** (4–25%)