RedLeaf

Isolation and Communication in a Safe Operating System

OSDI' 20

University of California, Irvine

VMware Research

TOC

1. Overview
2. Background
3. RedLeaf
4. Evaluation
5. Conclusion & Insight

1. Overview

Overview

- RedLeaf OS with novel **isolation** mechanism
  - NO costly hardware-based isolation
  - Relies on **type and memory safety of Rust**
  - **IDL** that supports cross-domain call proxying
  - (Engineering) POSIX-subset, NVMe, 10Gbps Intel ixgbe network

2. Background

2.1 如何评估一个安全的系统

- Domains: a unit of resources and info (本文中的主要分析对象)
- Domains can be **cleanly** terminated
- The **faults and crashes** in one domain do not affect other domains
- Shared objects cause many problems !

2.2 History of Isolation

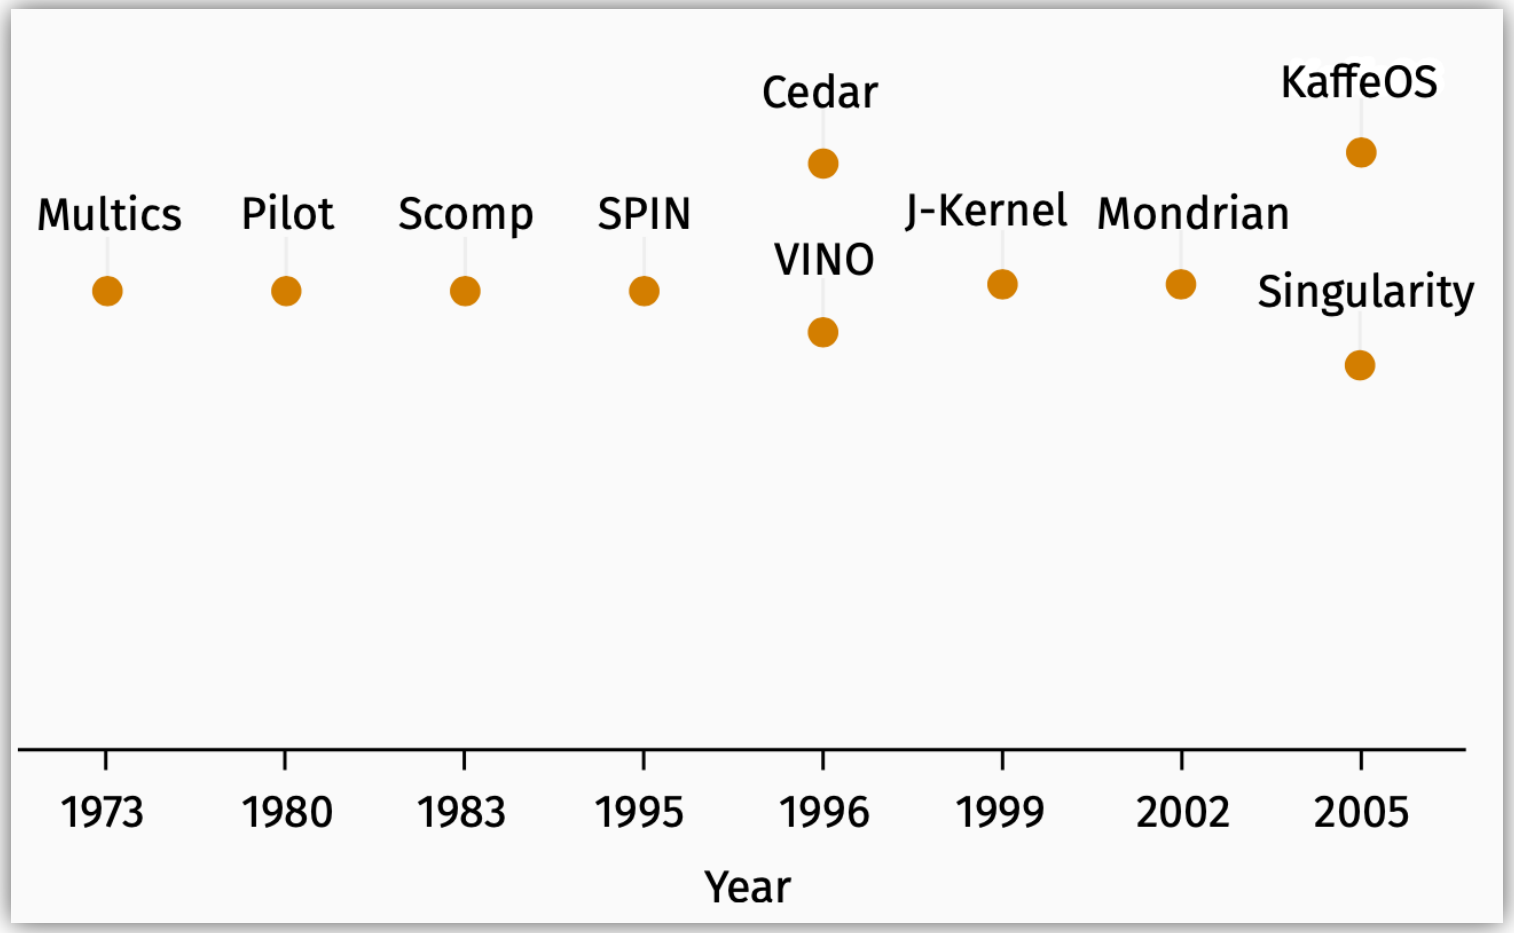

2.2 History of Isolation

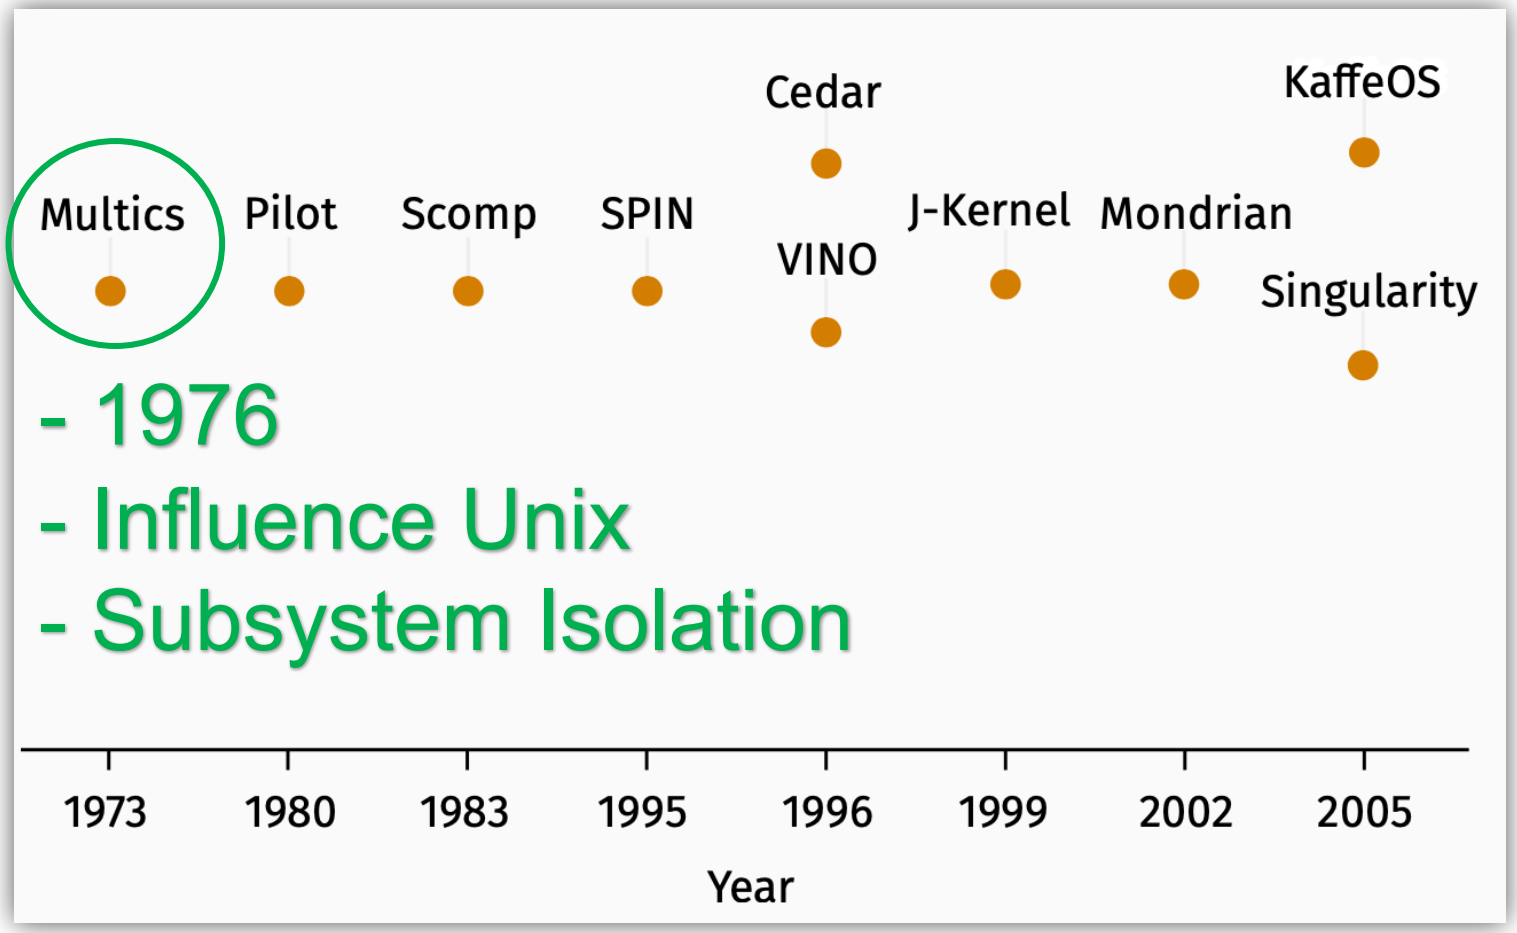

2.2 History of Isolation

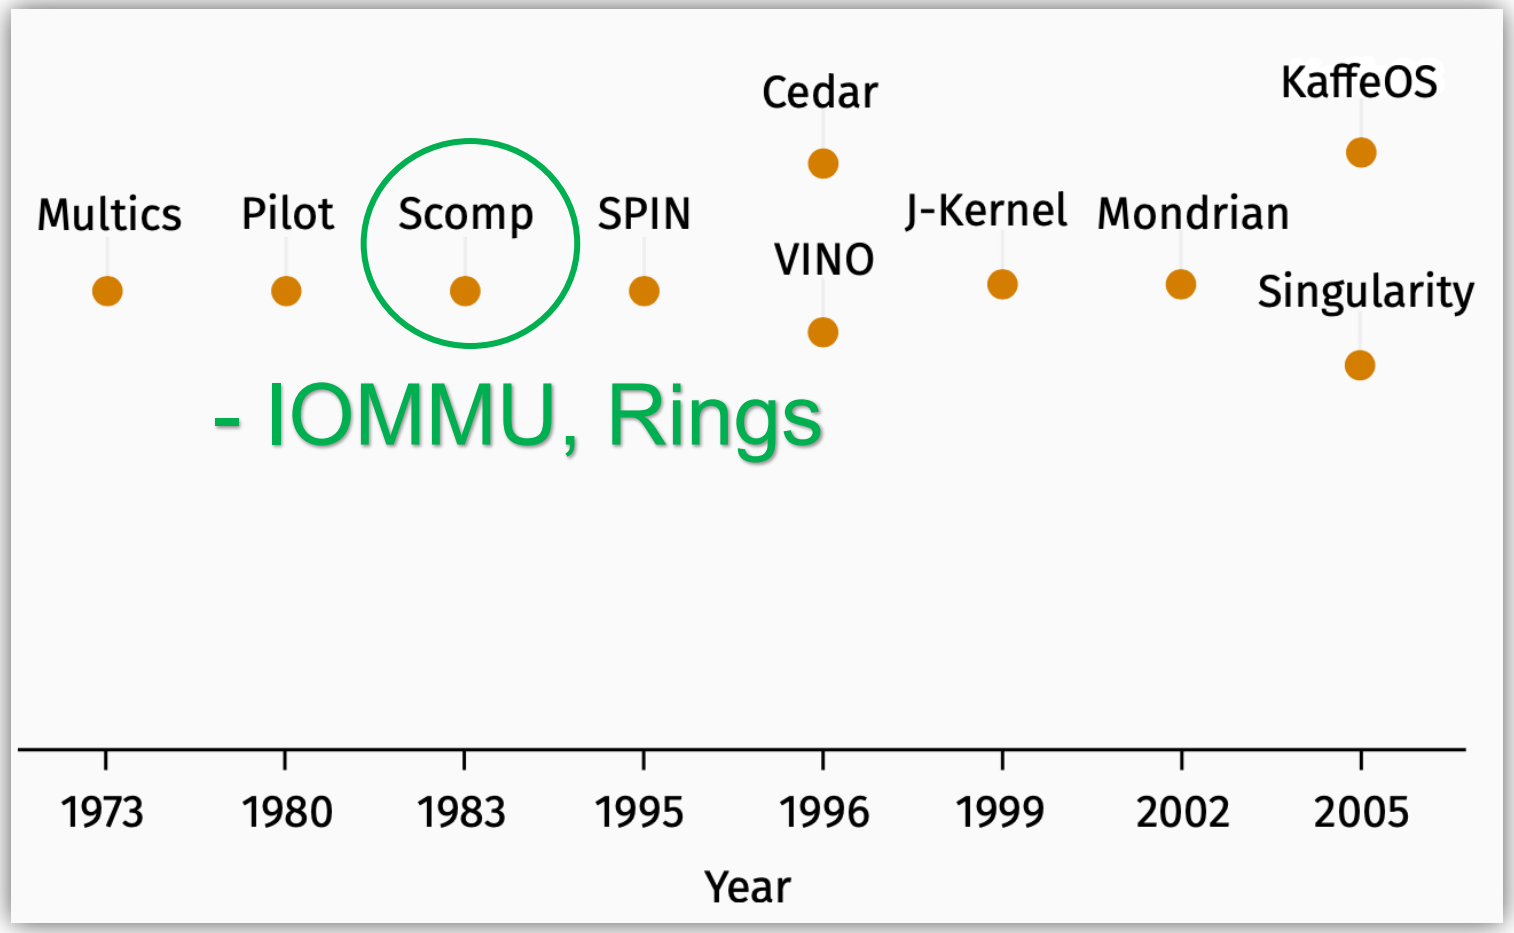

2.2 History of Isolation

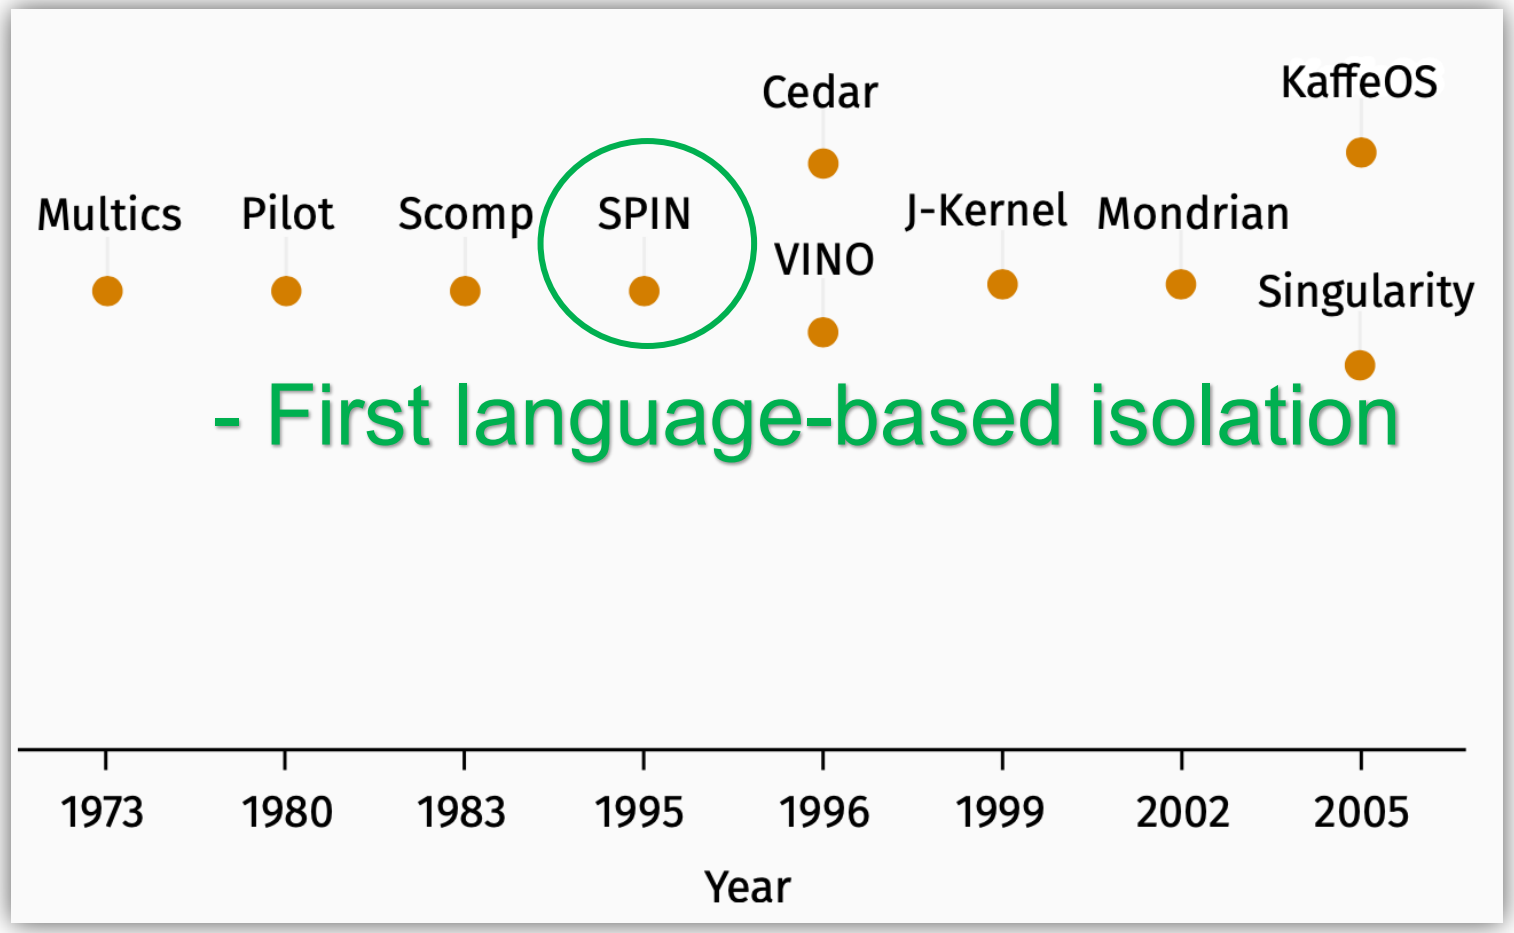

2.2 History of Isolation

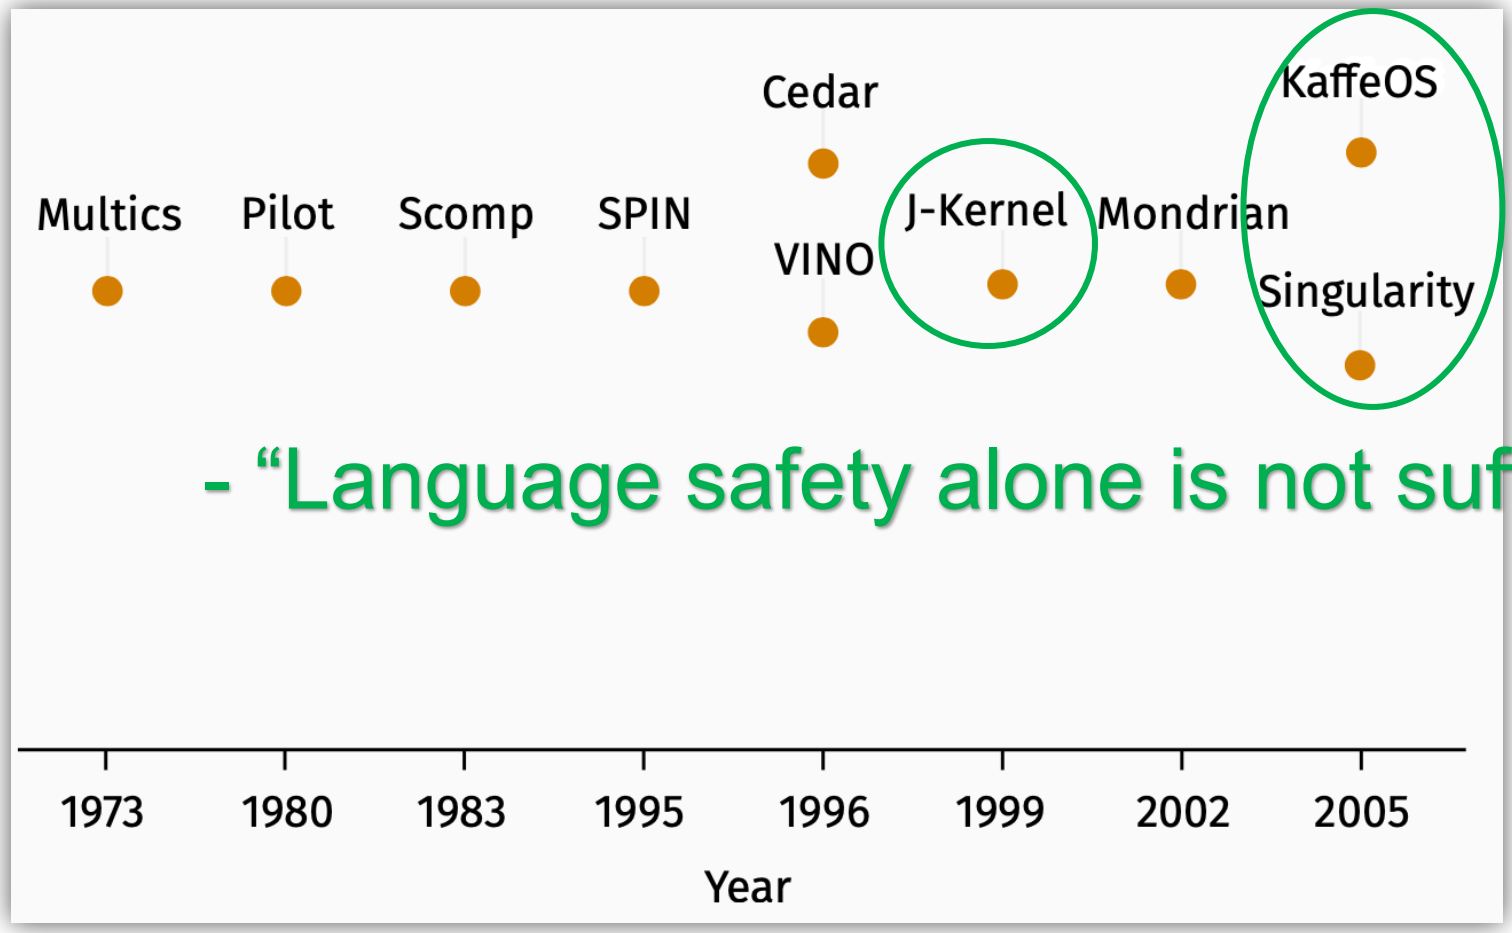

- “Language safety alone is not sufficient”

2.2 History of Isolation

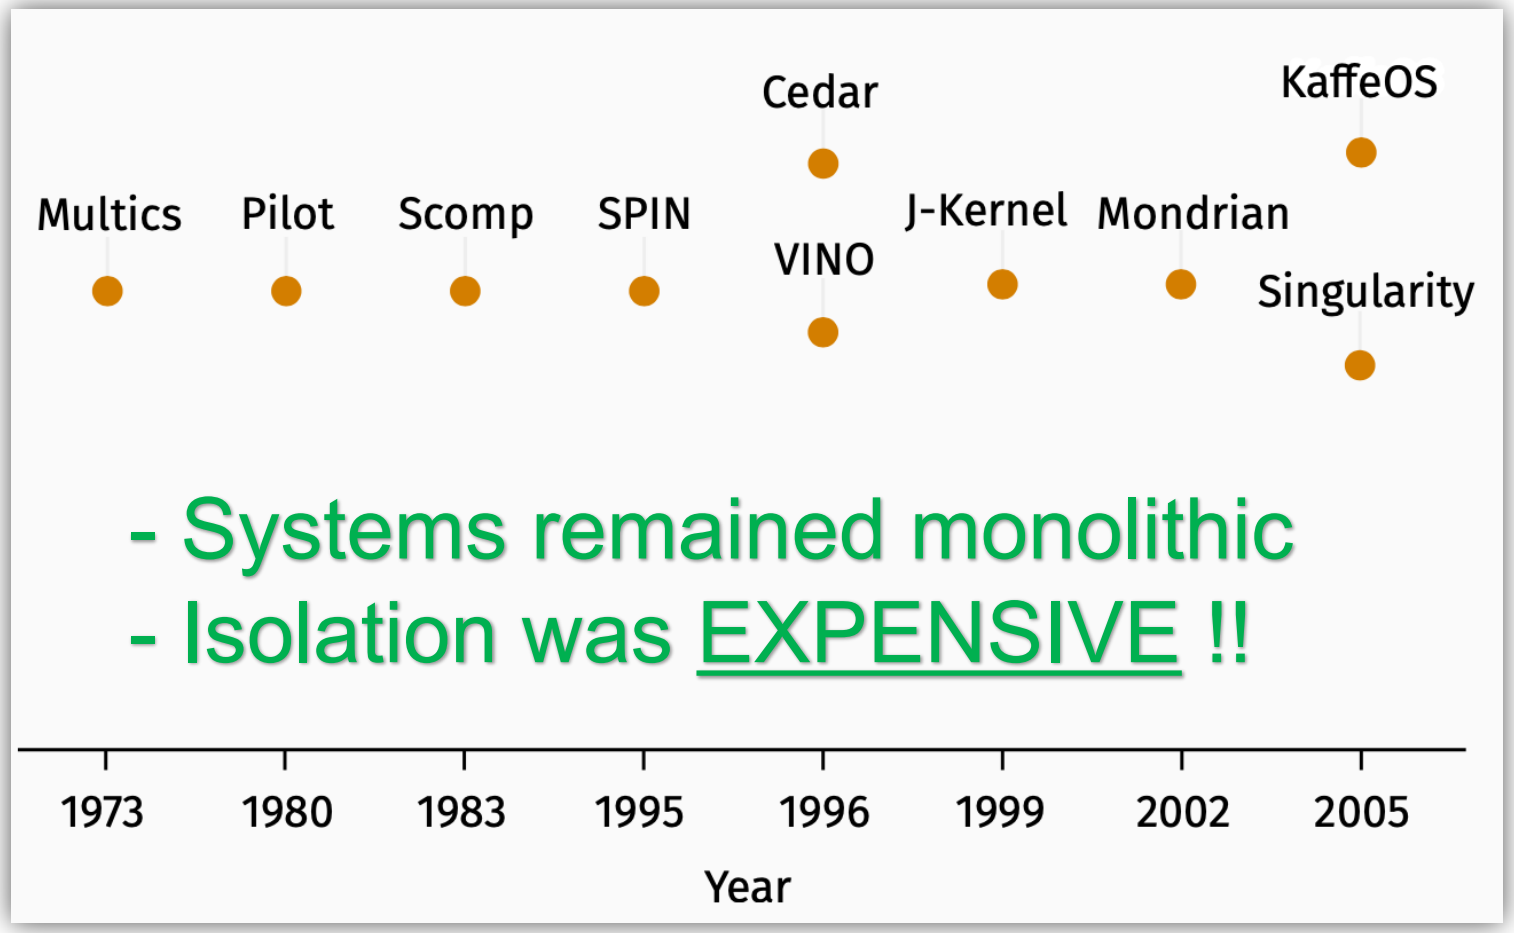

2.3 Isolation Mechanisms & Drawbacks

- Hardware Isolation & Latency
  - Segmentation (46 cycles)
  - Page table isolation (797 cycles)
  - VMFUNC (396 cycles)
  - Memory protection keys (20-26 cycles)
- Language based isolation
  - Compare drivers written (DPDK-style) in a safe high-level language (C, Rust, Go, C#, etc.)
  - Managed runtime and Garbage collection (20-50% overhead on a device-driver workload)

2.4 Traditional Safe Languages

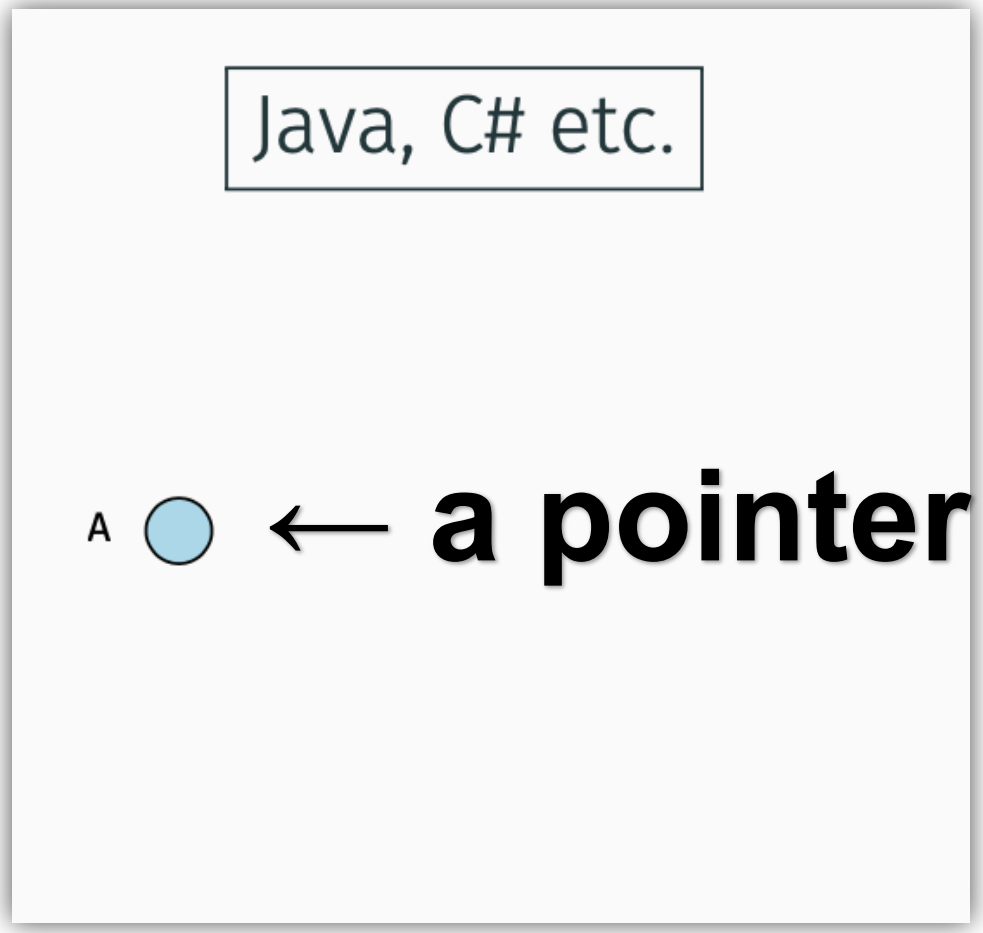

2.4 Traditional Safe Languages

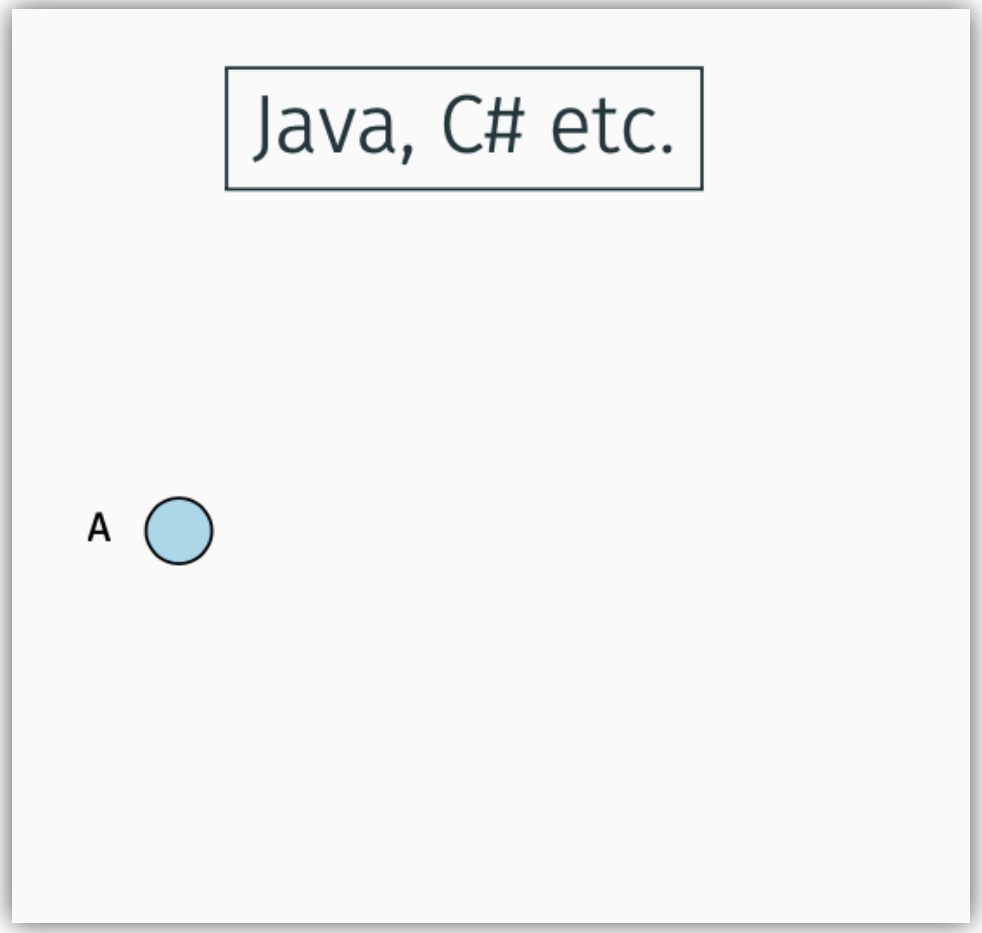

2.4 Traditional Safe Languages

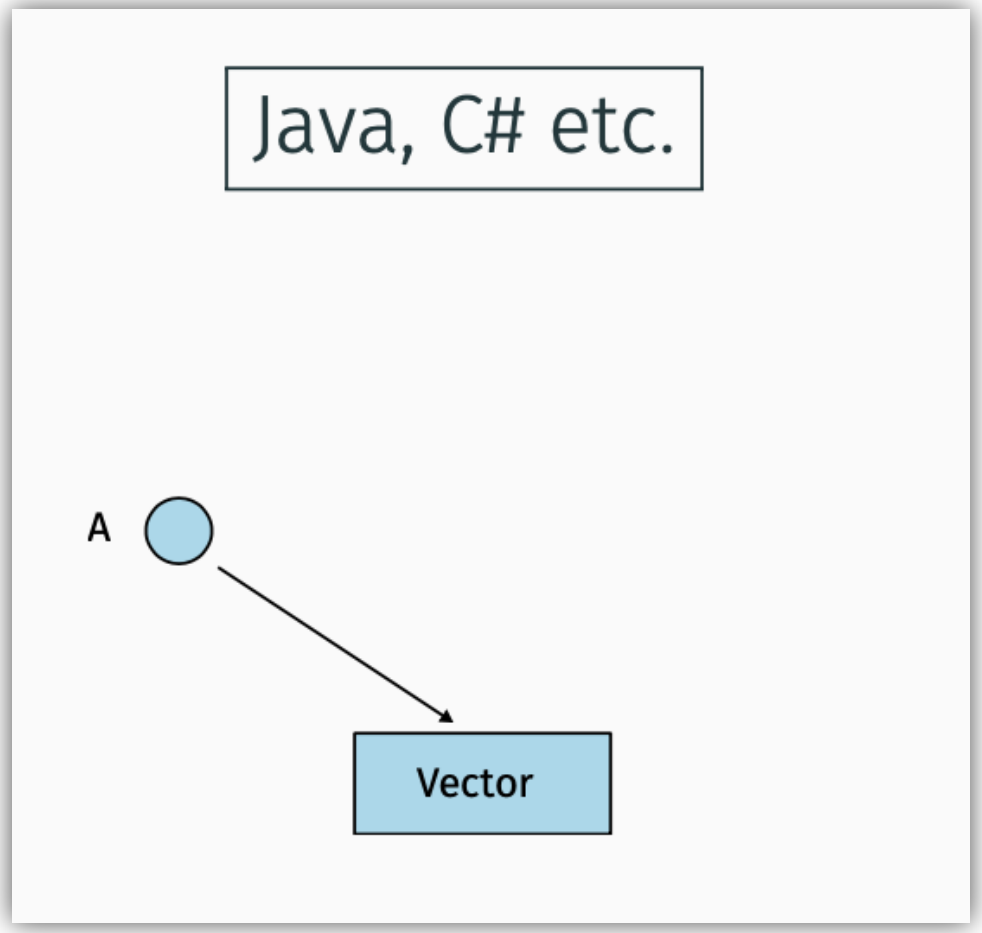

2.4 Traditional Safe Languages

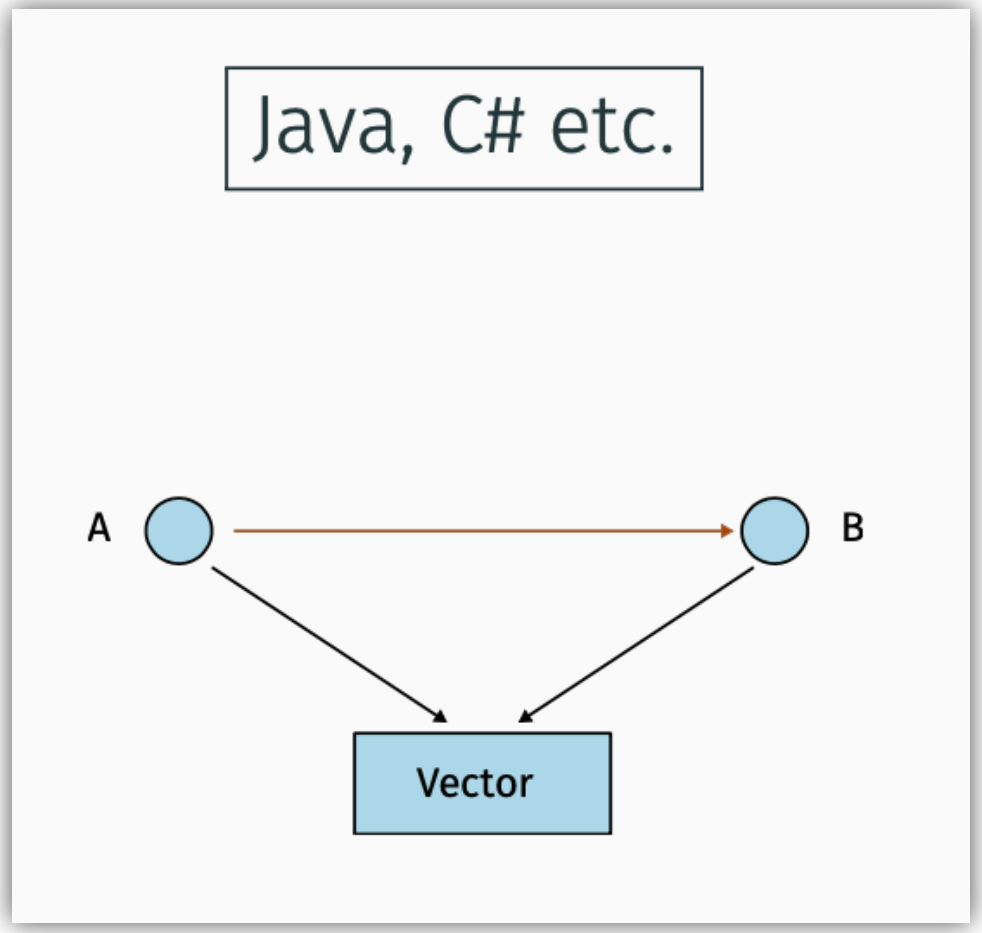

2.4 Traditional Safe Languages

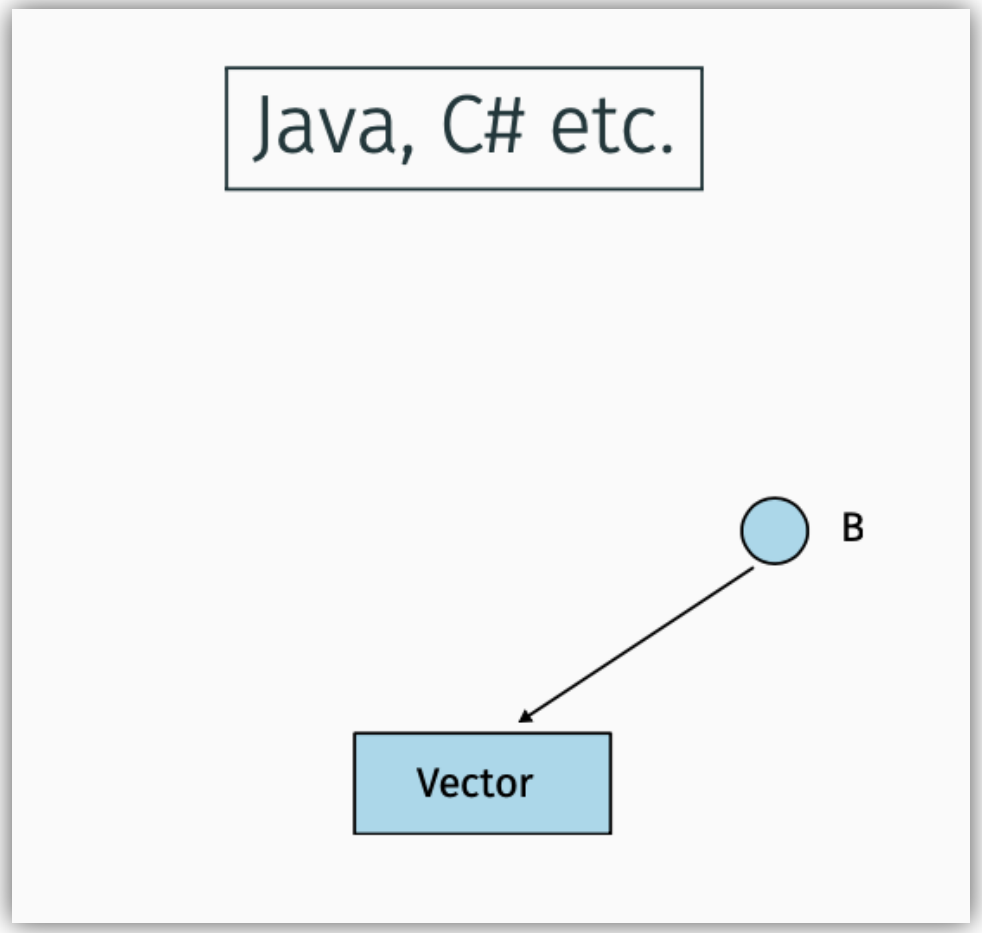

2.4 Traditional Safe Languages

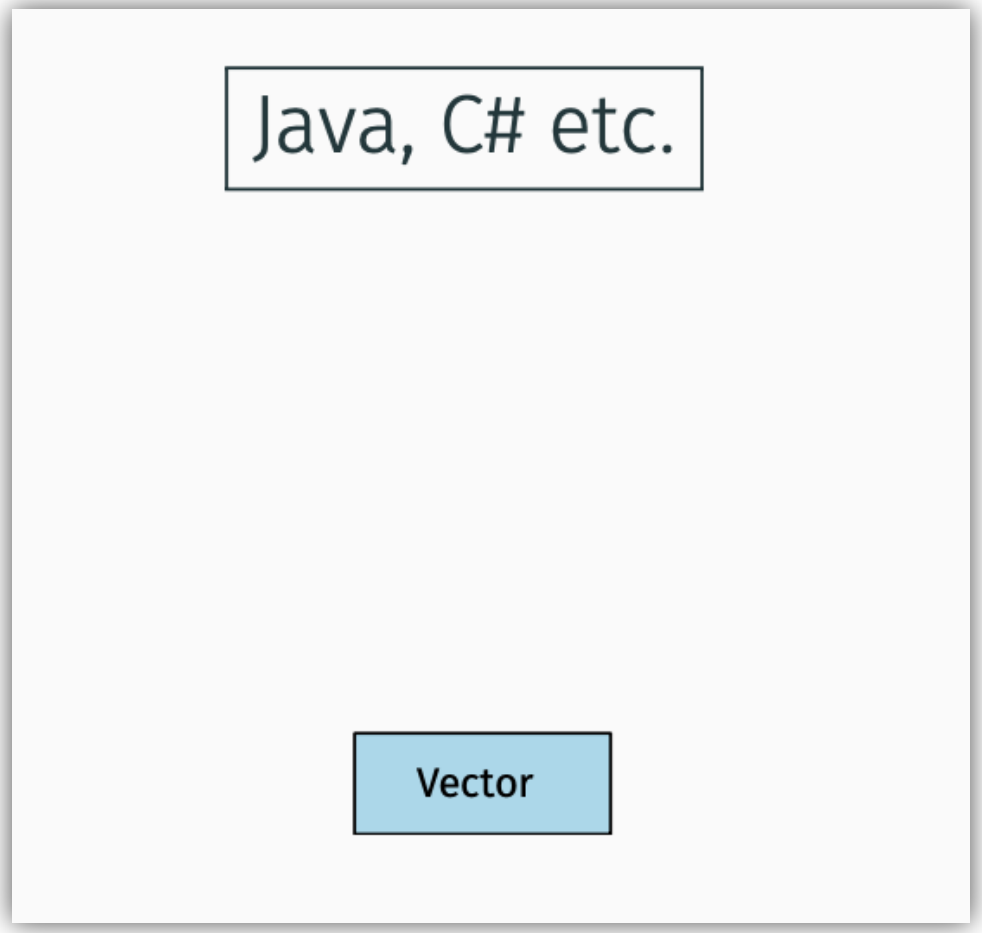

2.4 Traditional Safe Languages

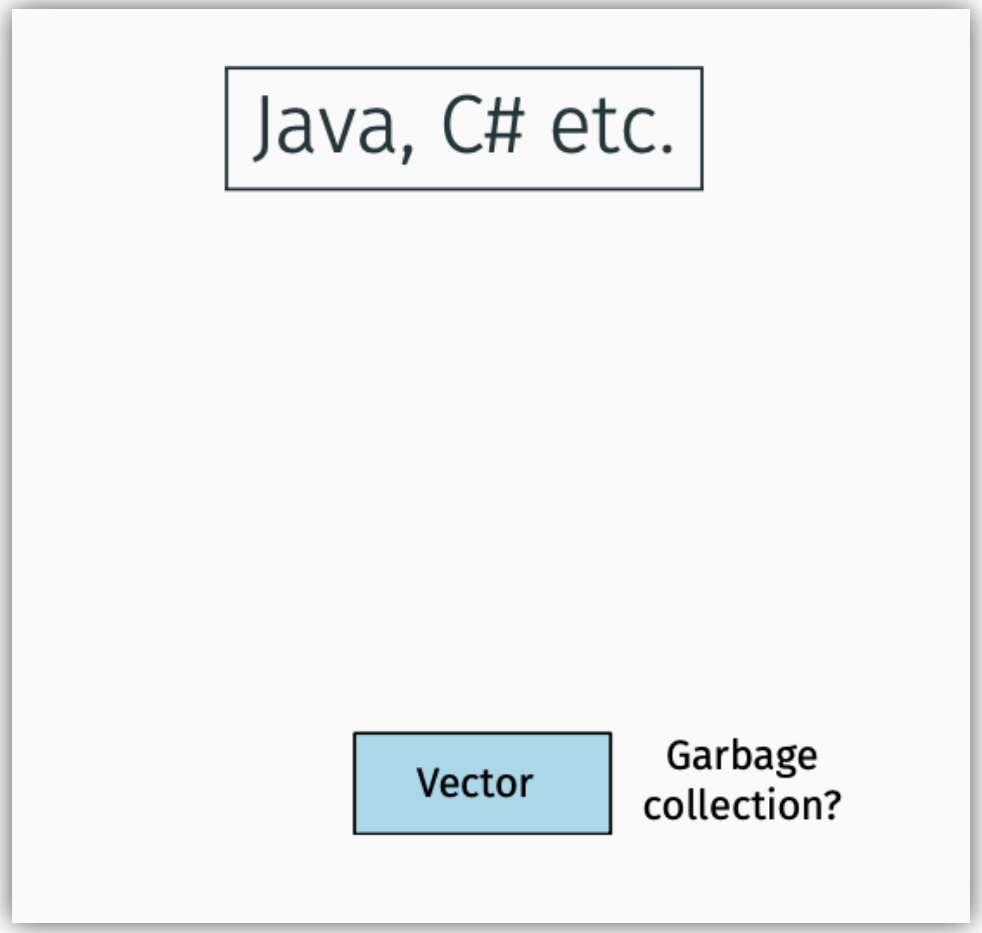

2.4 Traditional Safe Languages

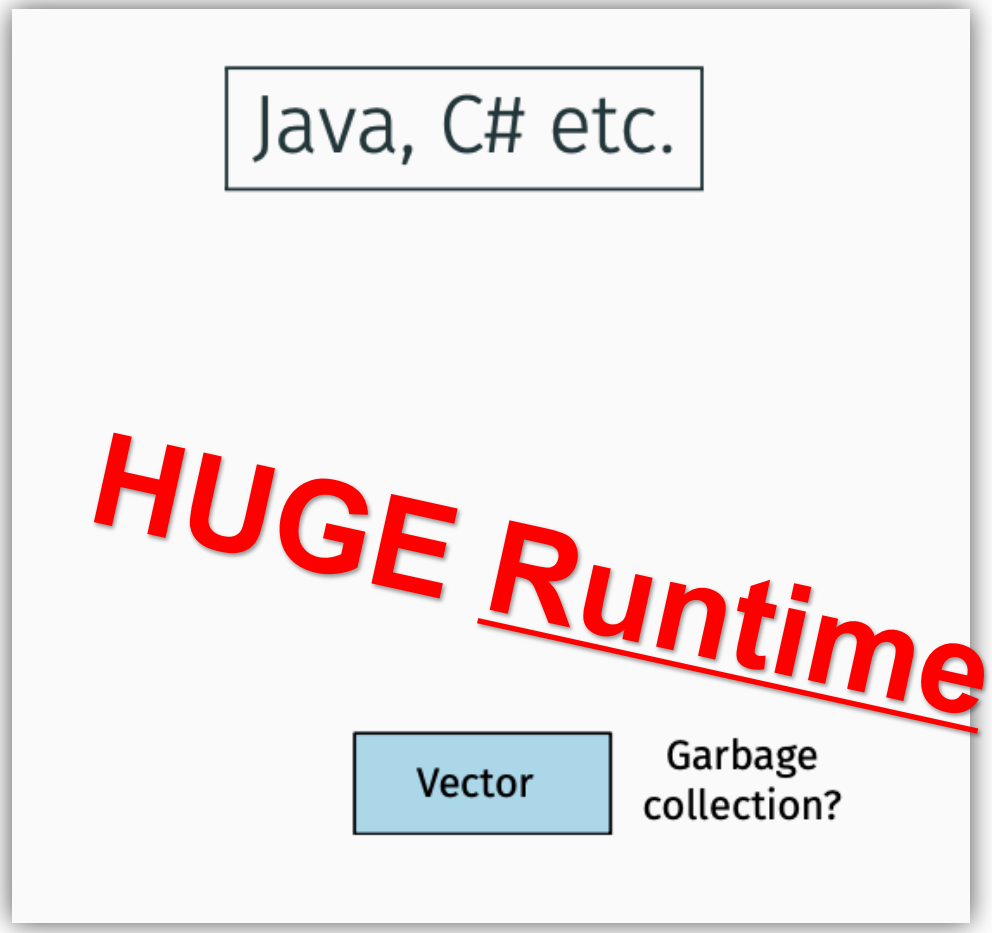

2.5 Rust

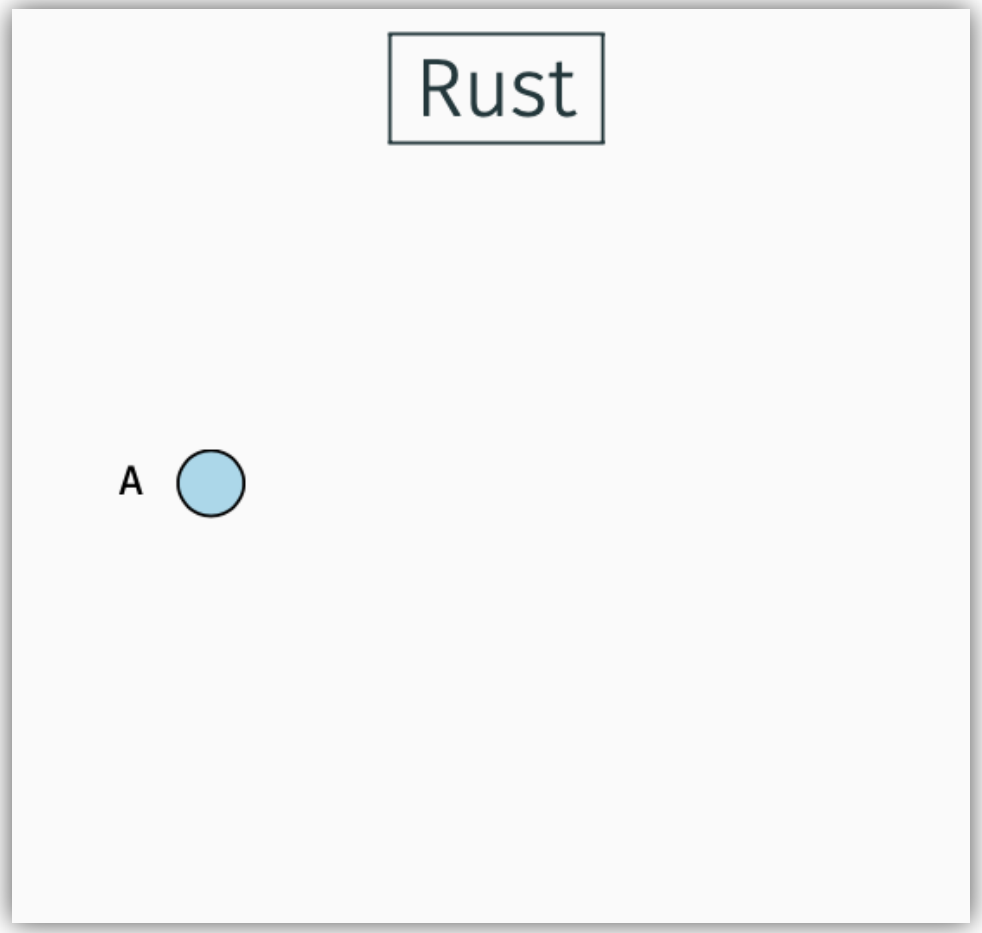

2.5 Rust

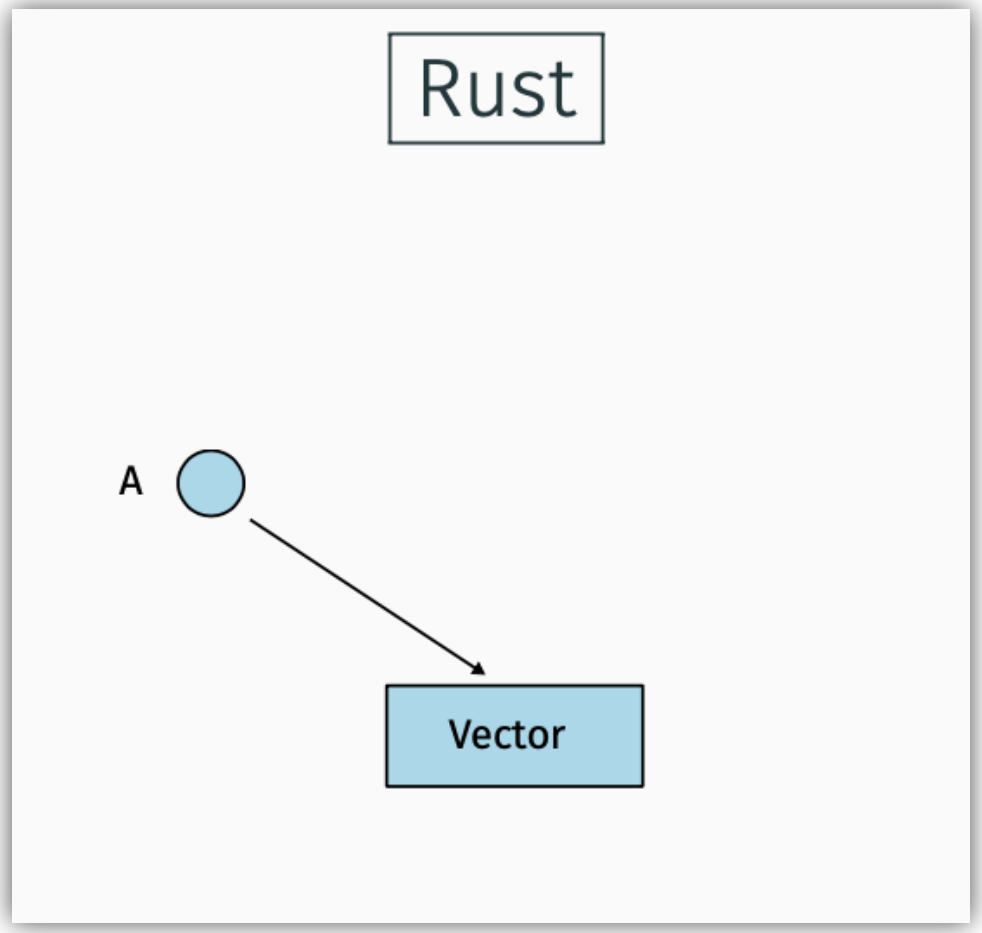

2.5 Rust

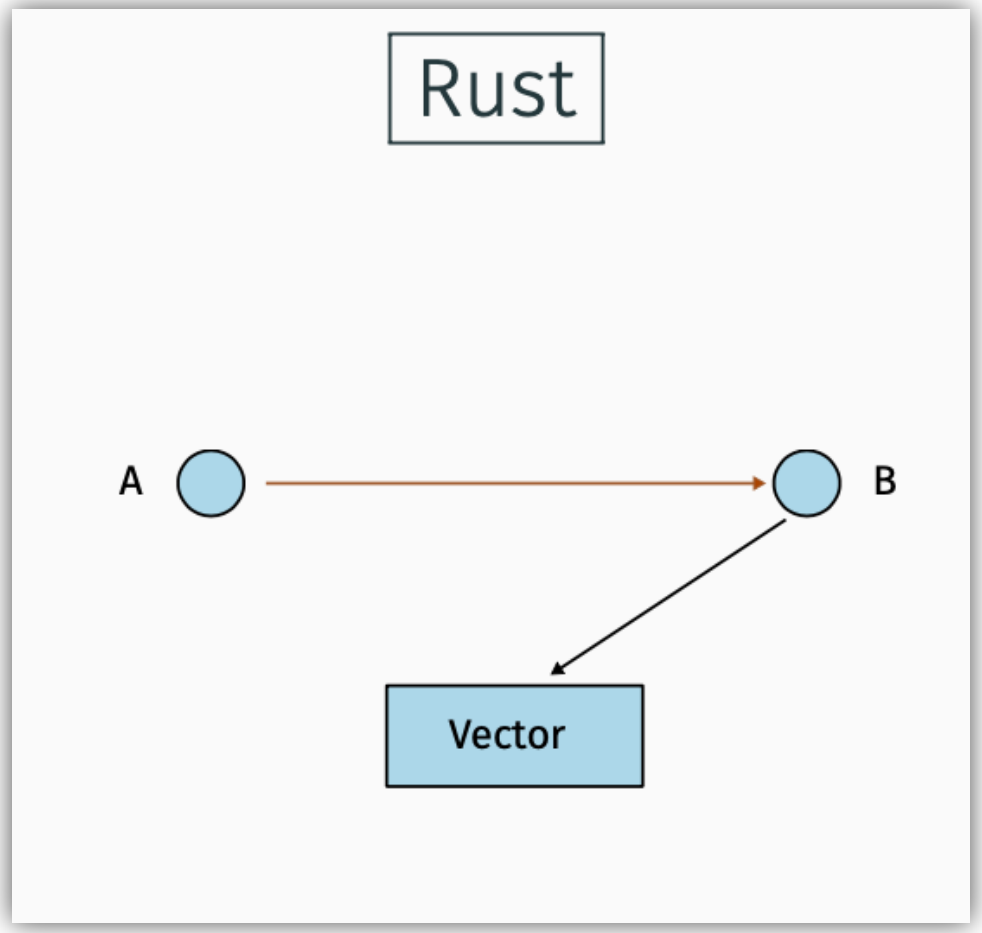

2.5 Rust

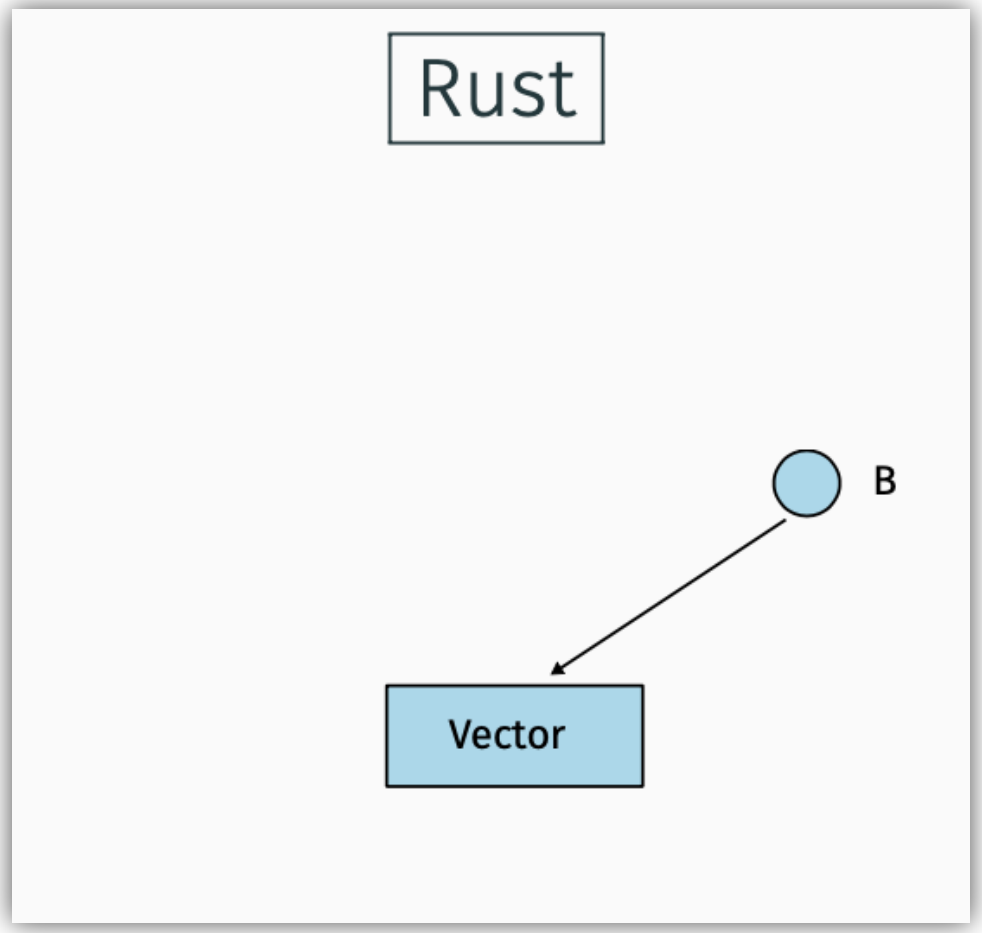

2.5 Rust

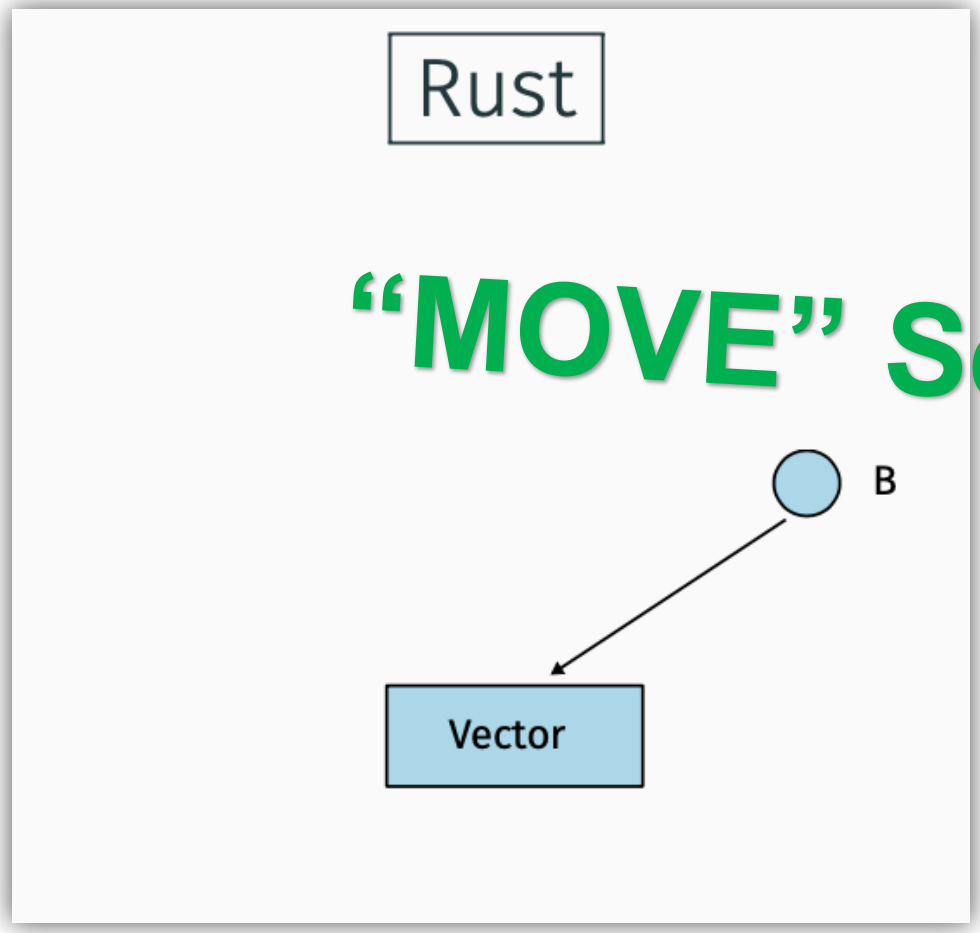

2.5 Rust

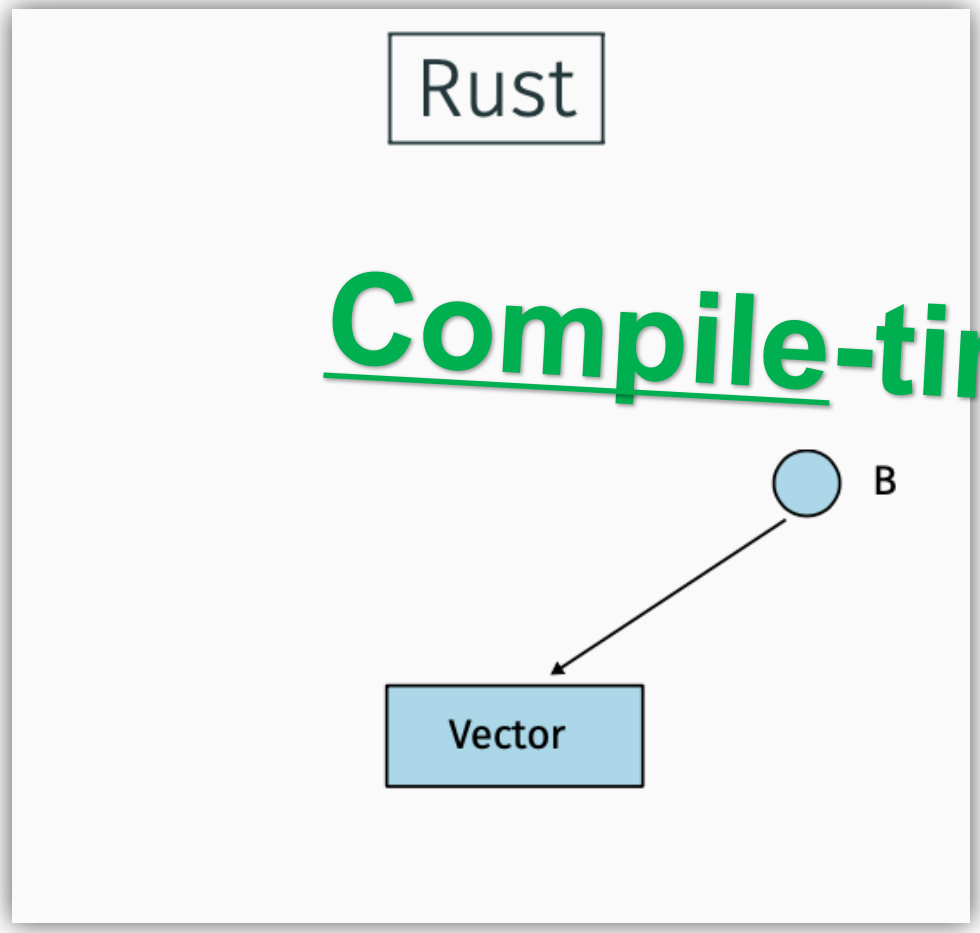

2.5 Rust

- Mostly use Rust as a drop-in replacement for C
- Numerous possibilities
  - Fault Isolation
  - Transparent device-driver recovery
  - Safe Kernel extensions
  - Fine-grained capability-based access control etc.

2.6 Fault Isolation in Language-based Systems

- Fault Isolation as is mentioned before:
  - Domains can be **cleanly** terminated
  - The **faults and crashes** in one domain do not affect other domains

2.6 Fault Isolation in Language-based Systems

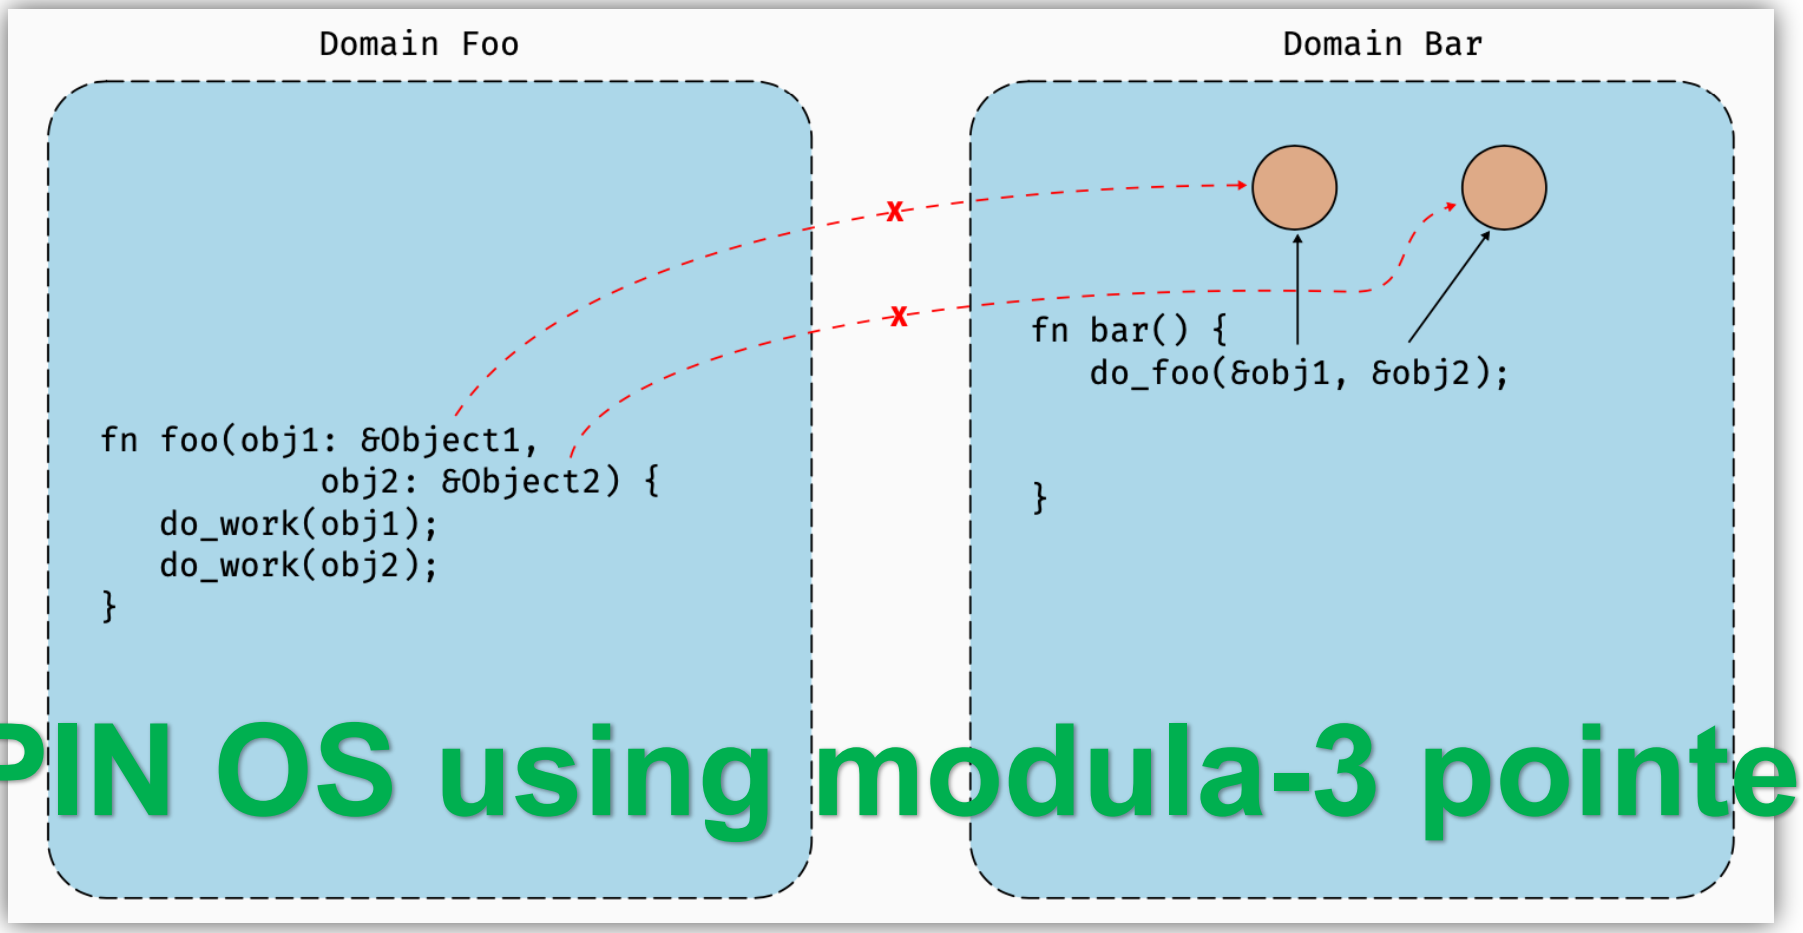

SPIN OS using modula-3 pointers

2.6 Fault Isolation in Language-based Systems

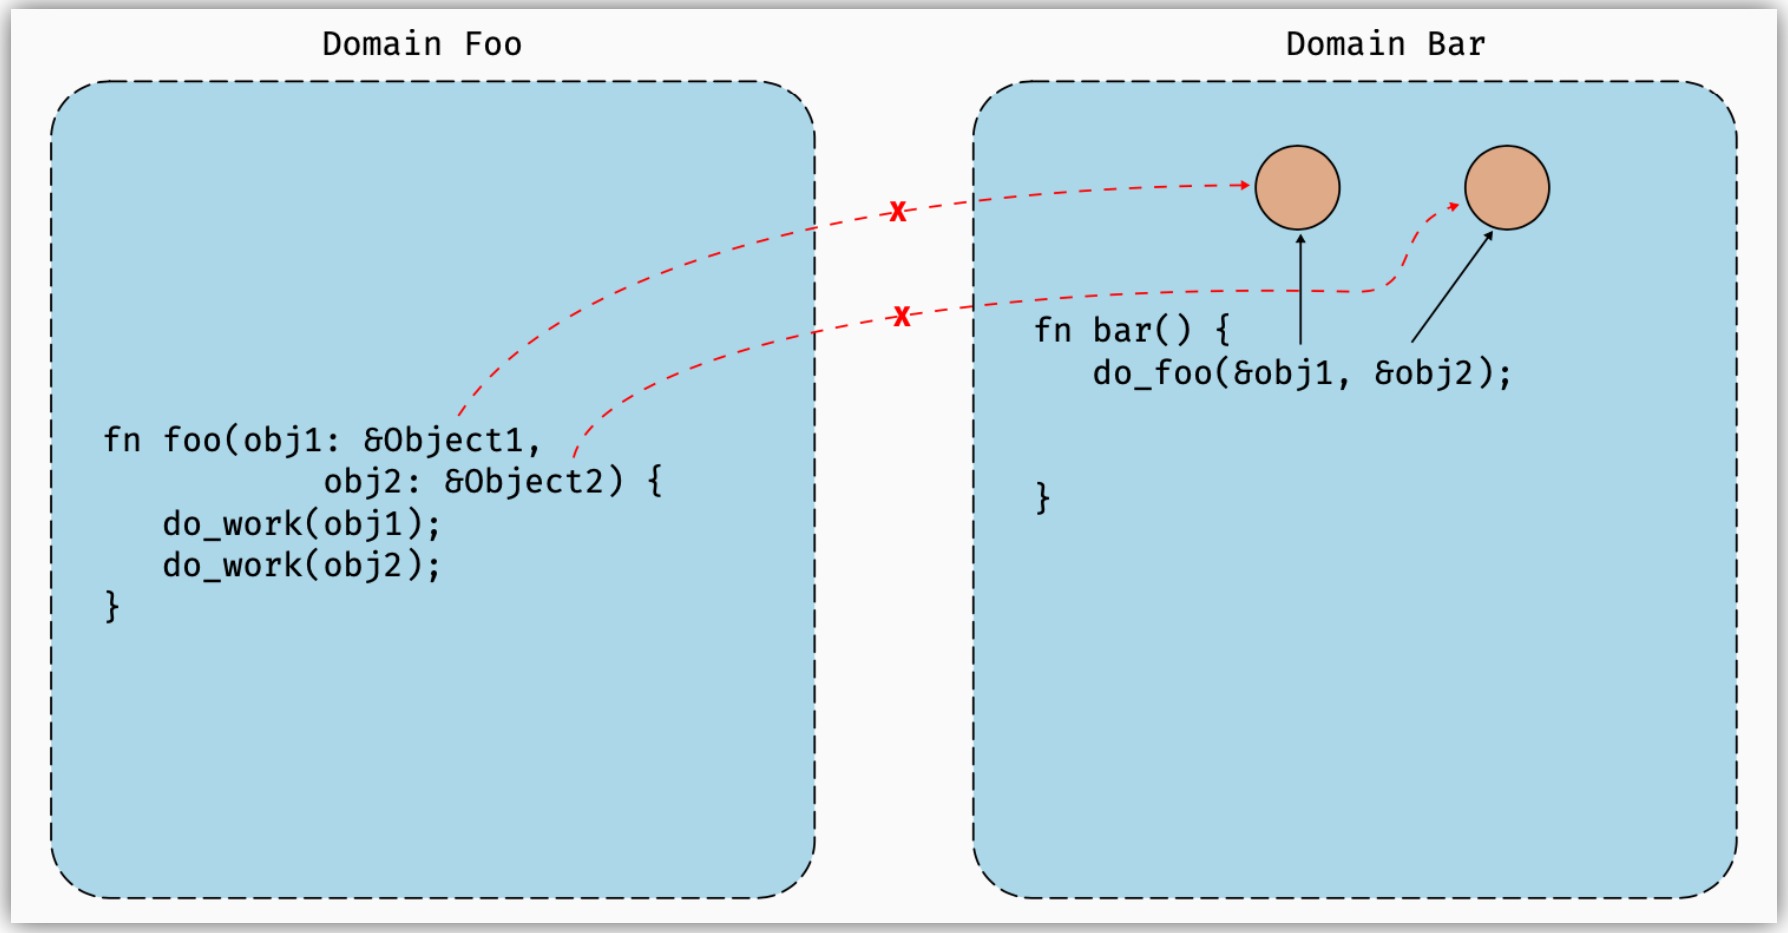

2.6 Fault Isolation in Language-based Systems

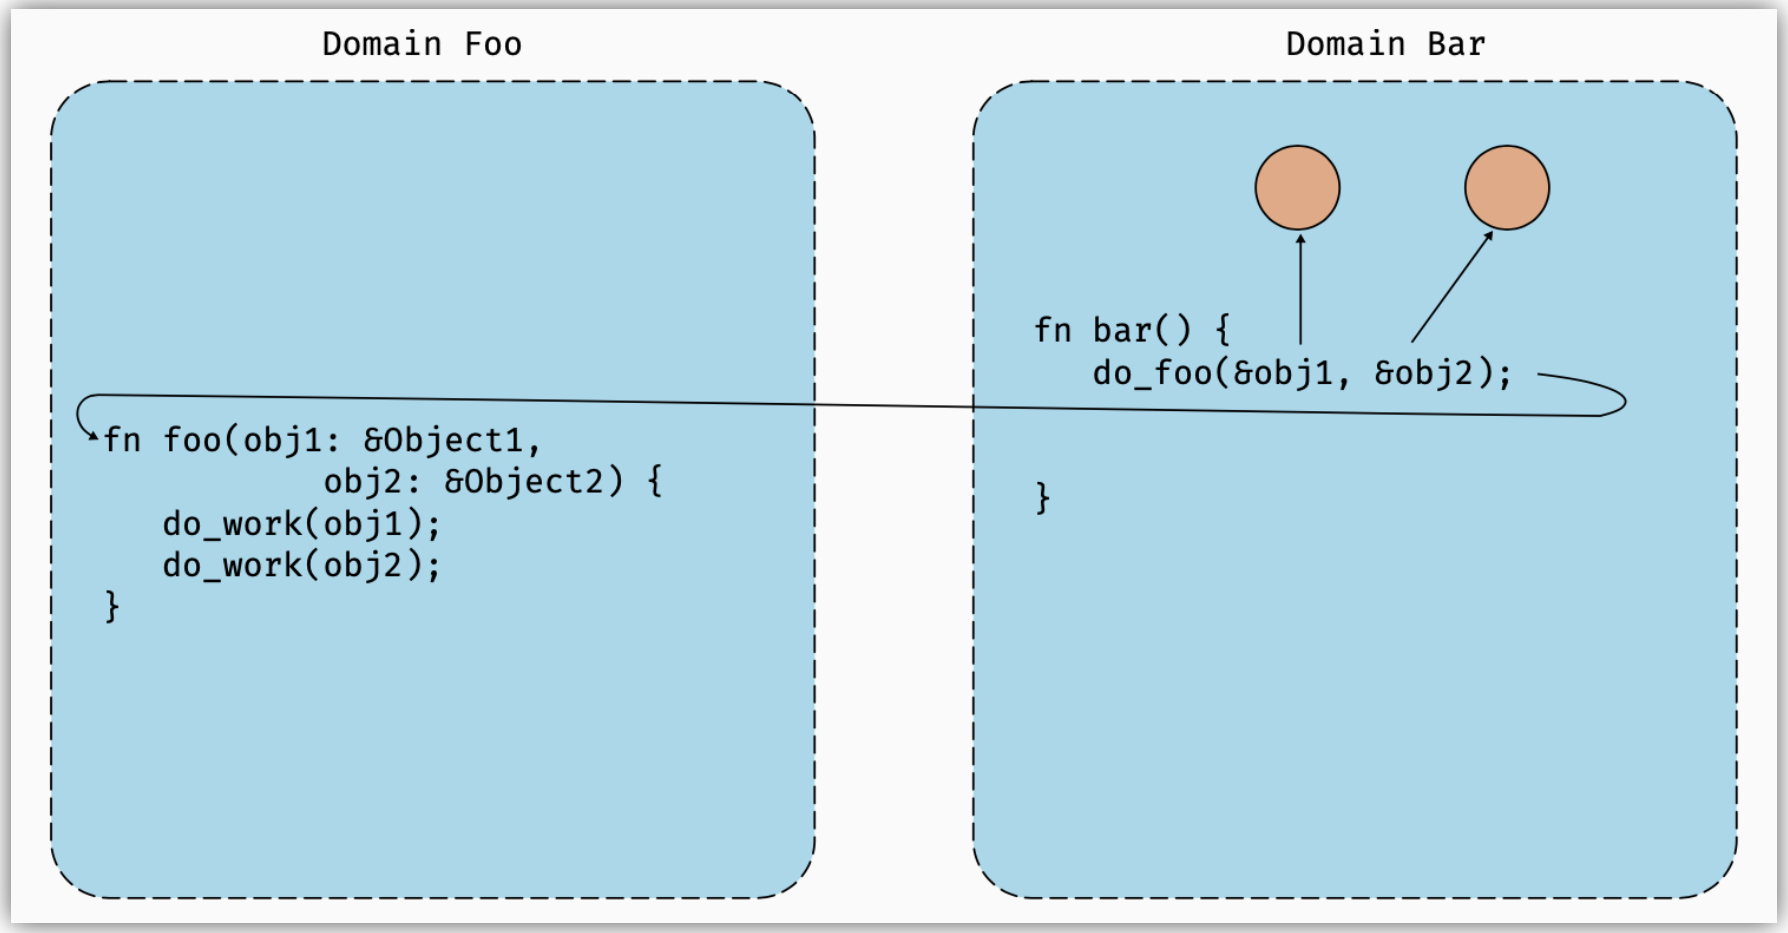

2.6 Fault Isolation in Language-based Systems

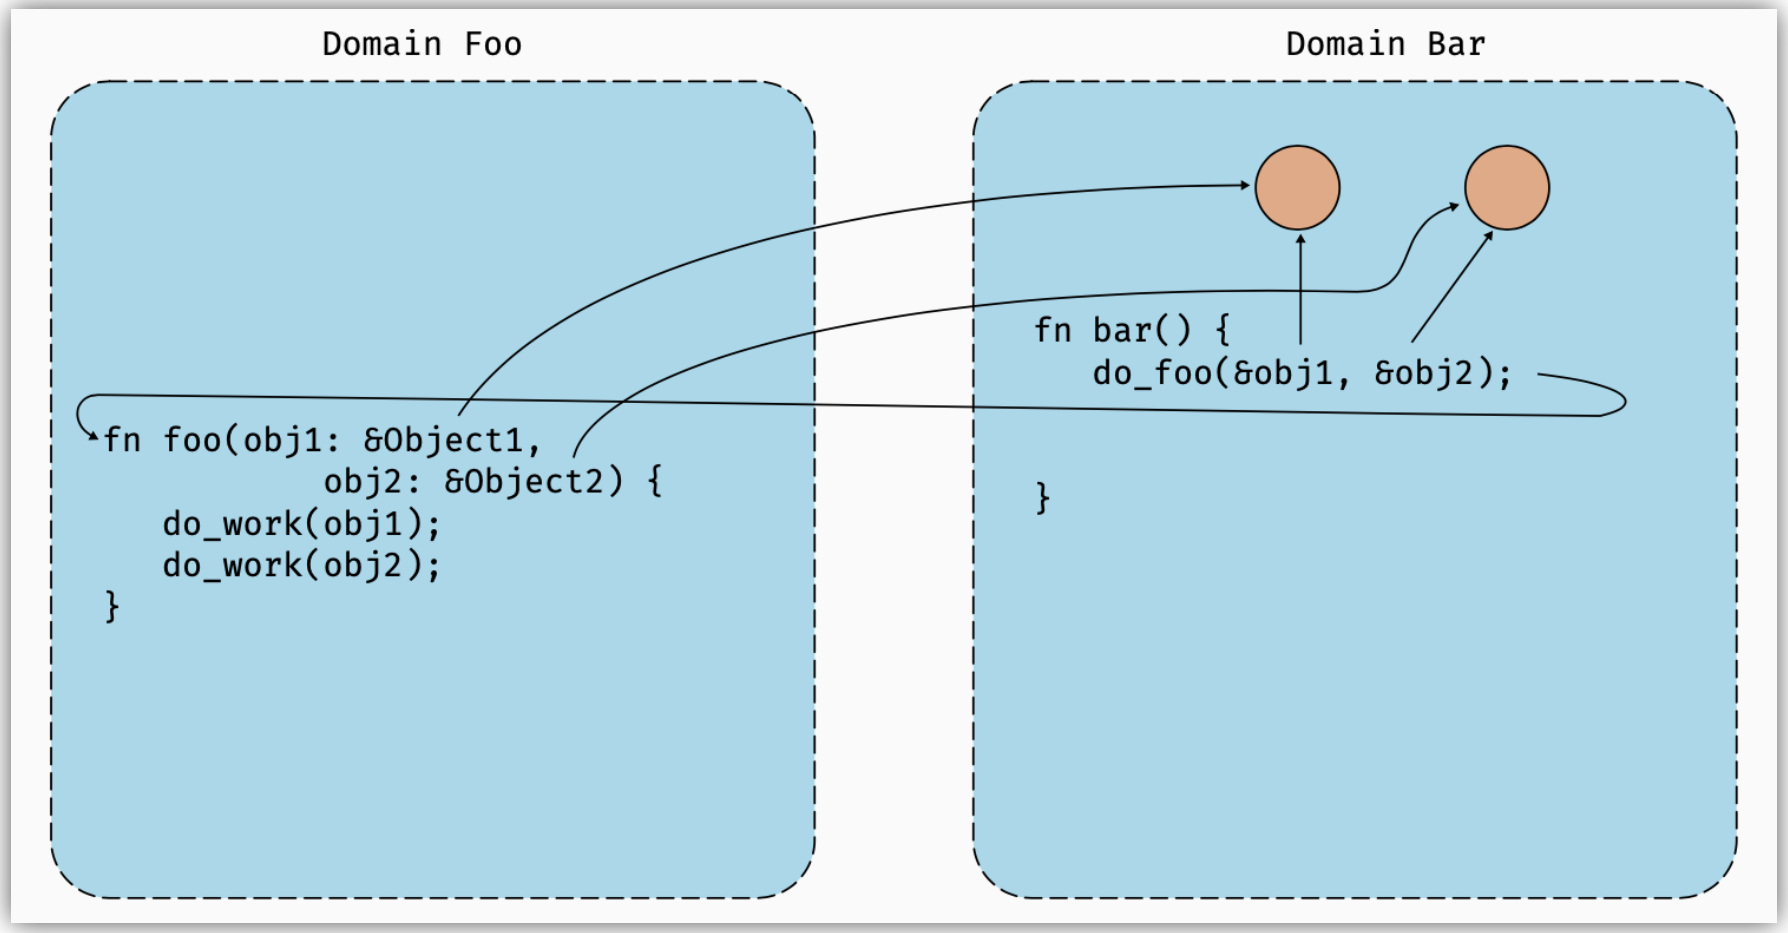

2.6 Fault Isolation in Language-based Systems

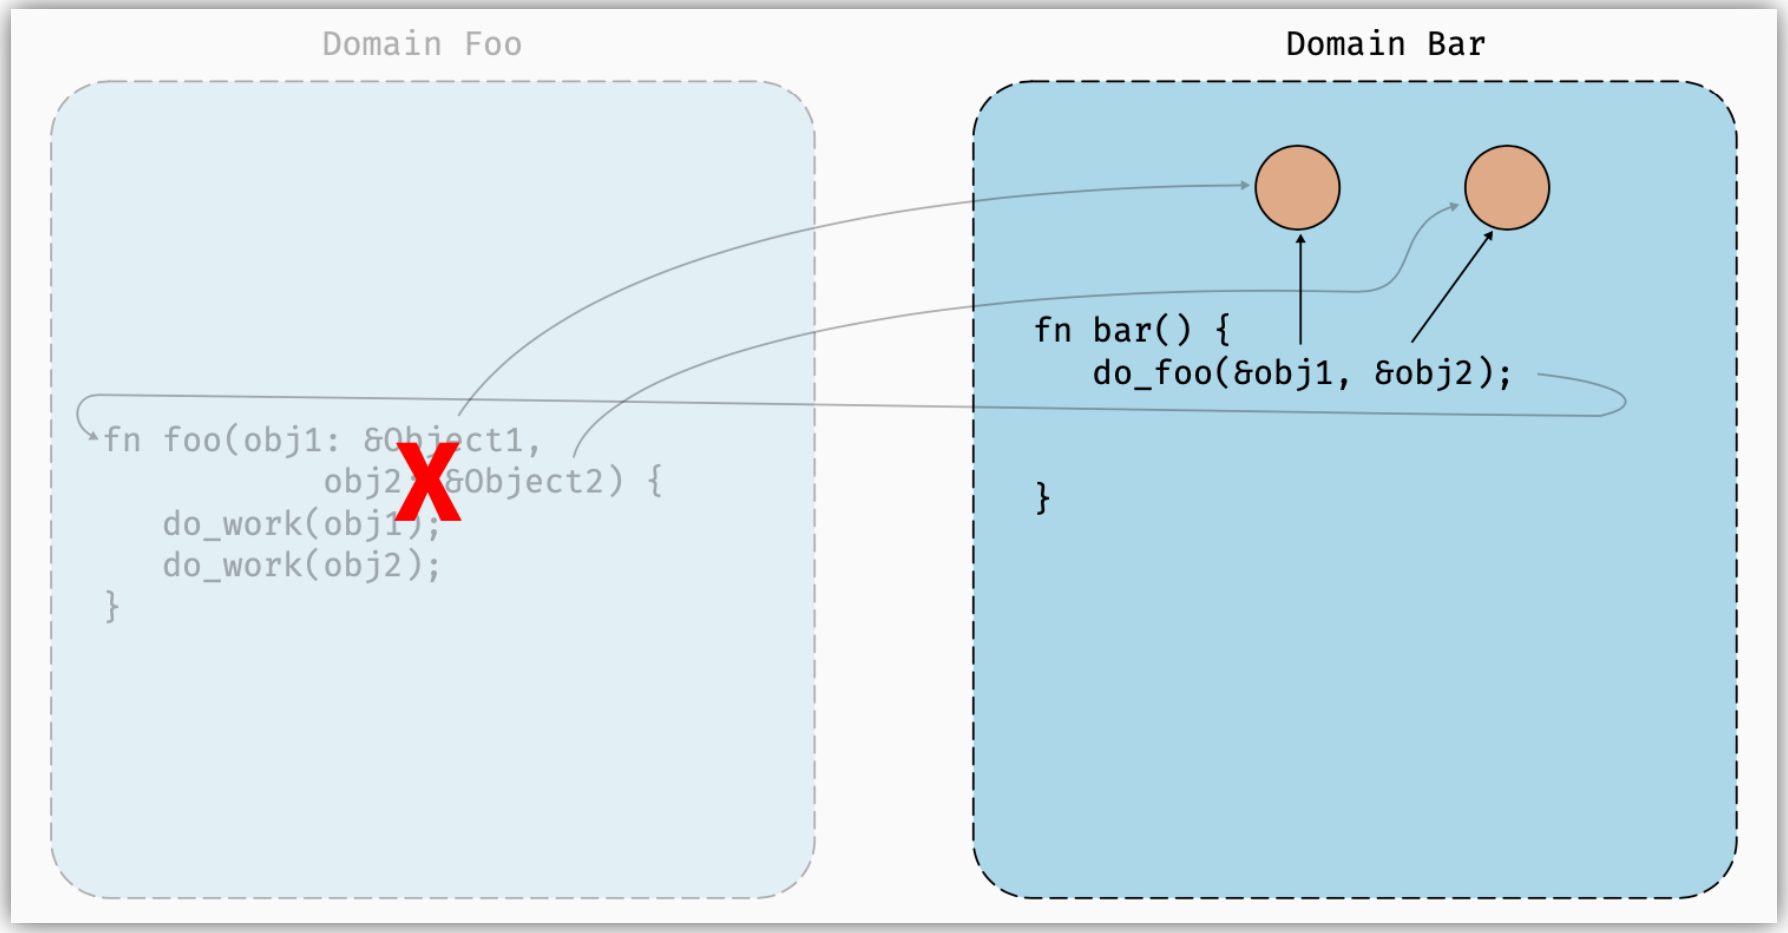

2.6 Fault Isolation in Language-based Systems

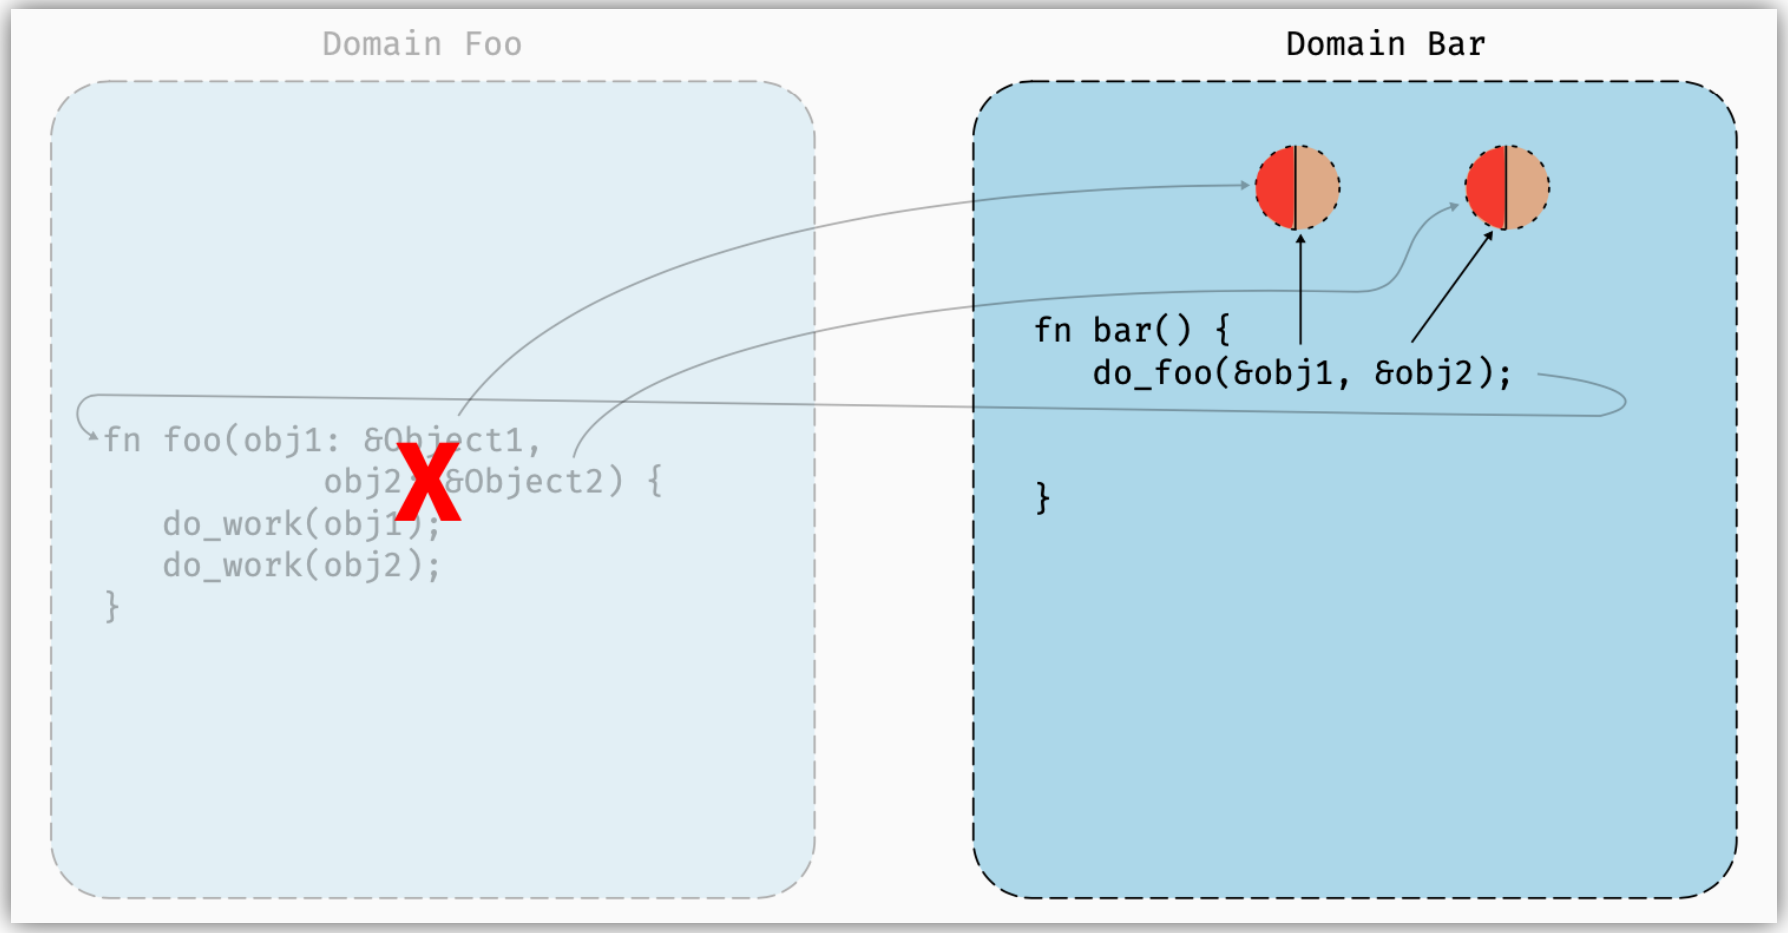

2.6 Fault Isolation in Language-based Systems

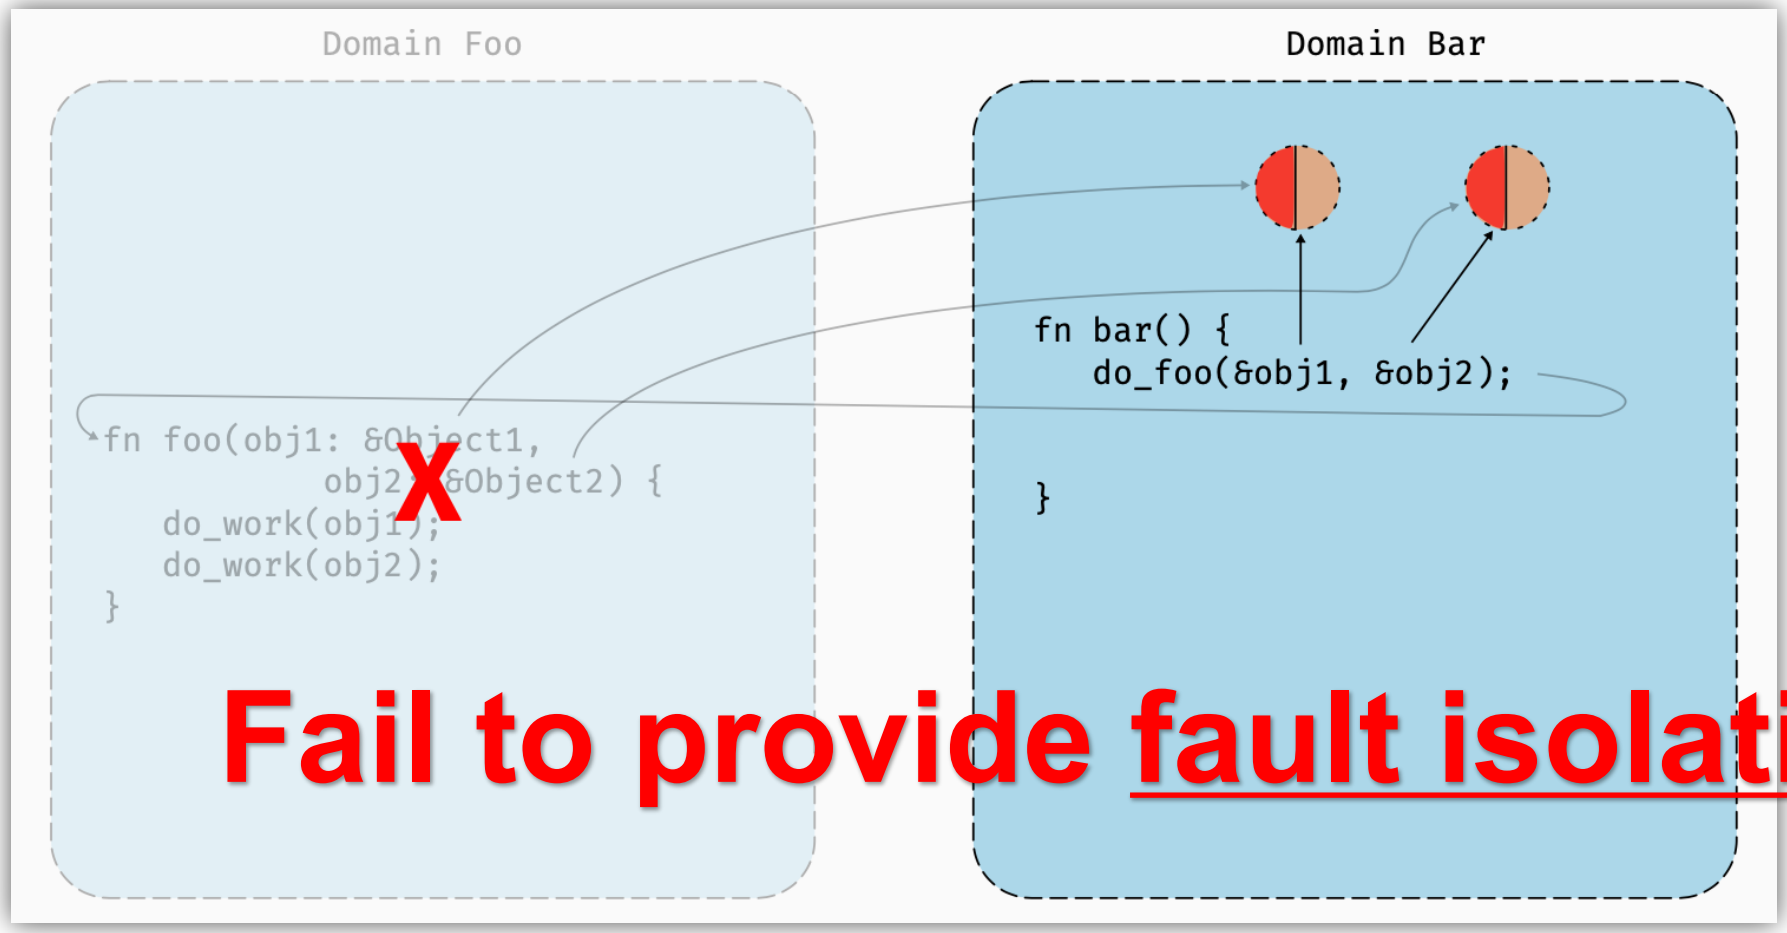

2.6 Fault Isolation in Language-based Systems: Deep Copy

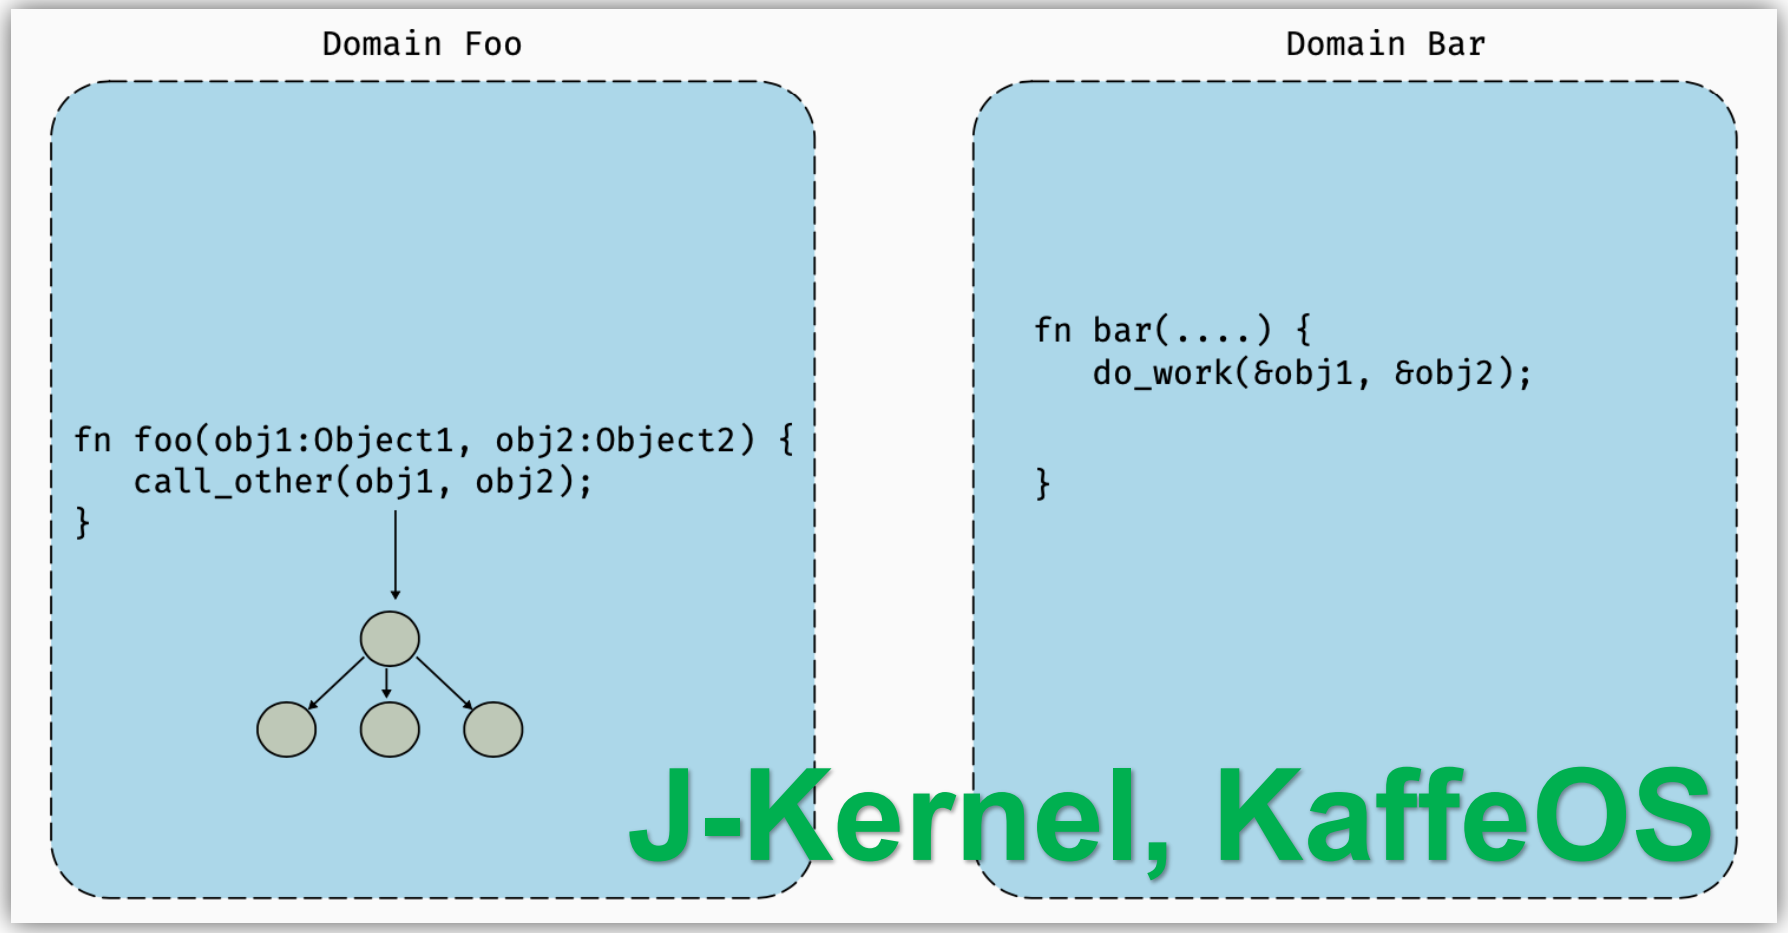

2.6 Fault Isolation in Language-based Systems: Deep Copy

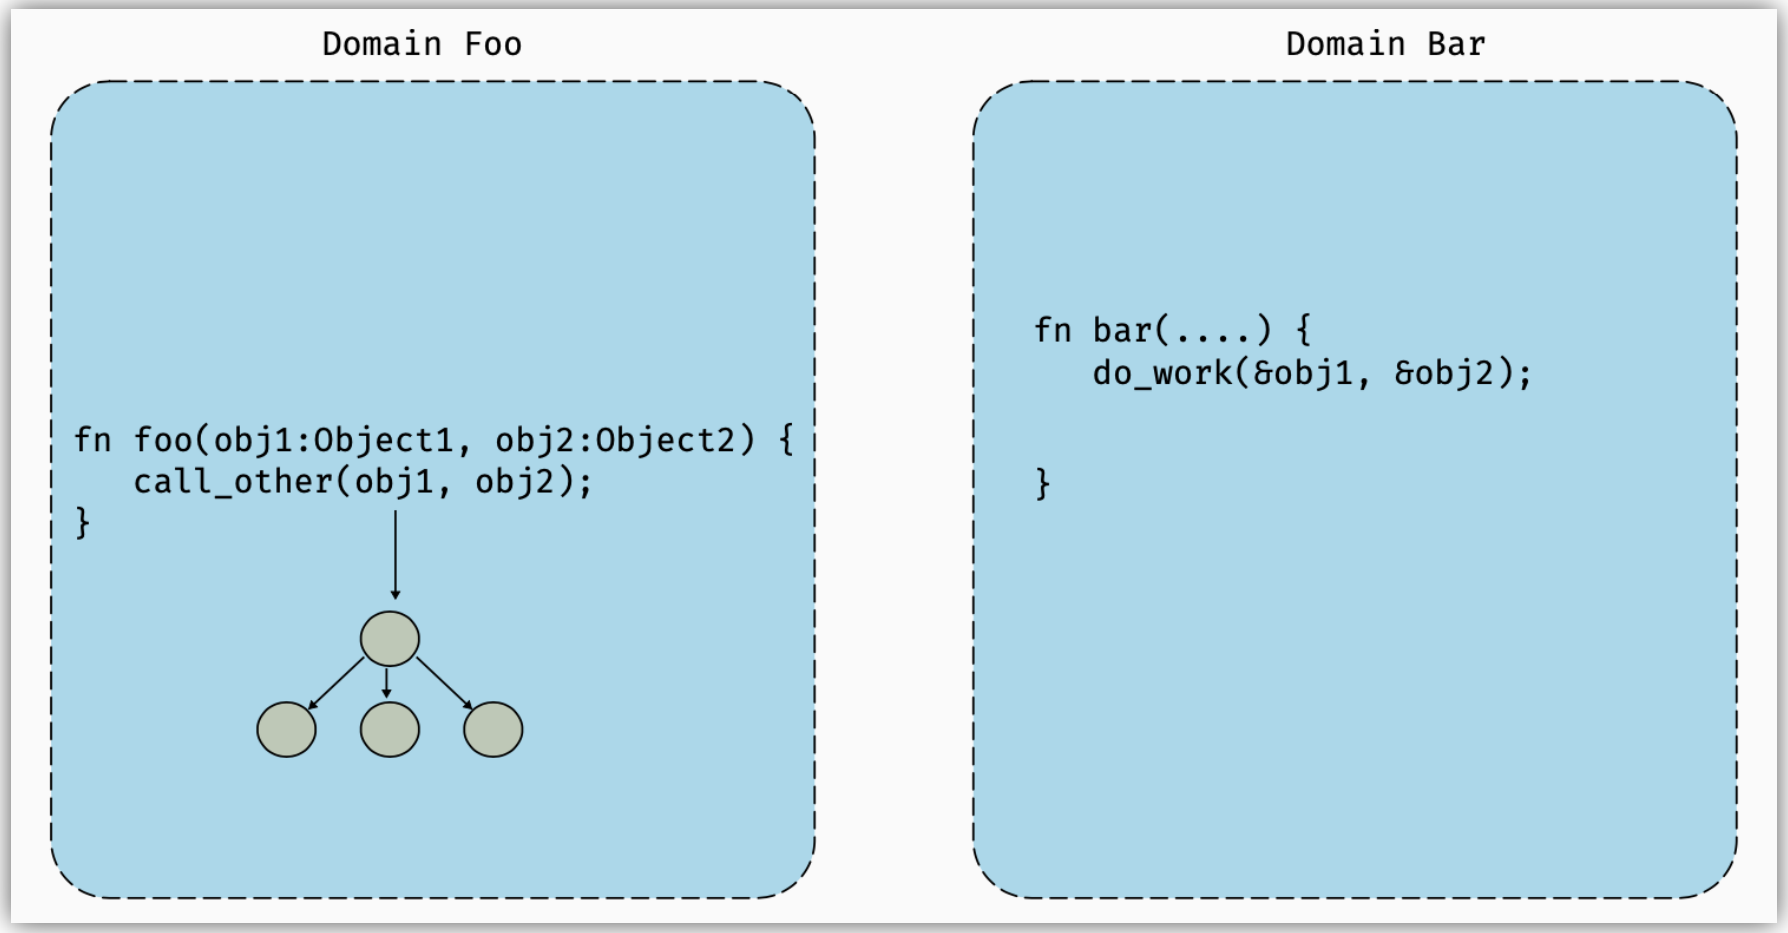

2.6 Fault Isolation in Language-based Systems: Deep Copy

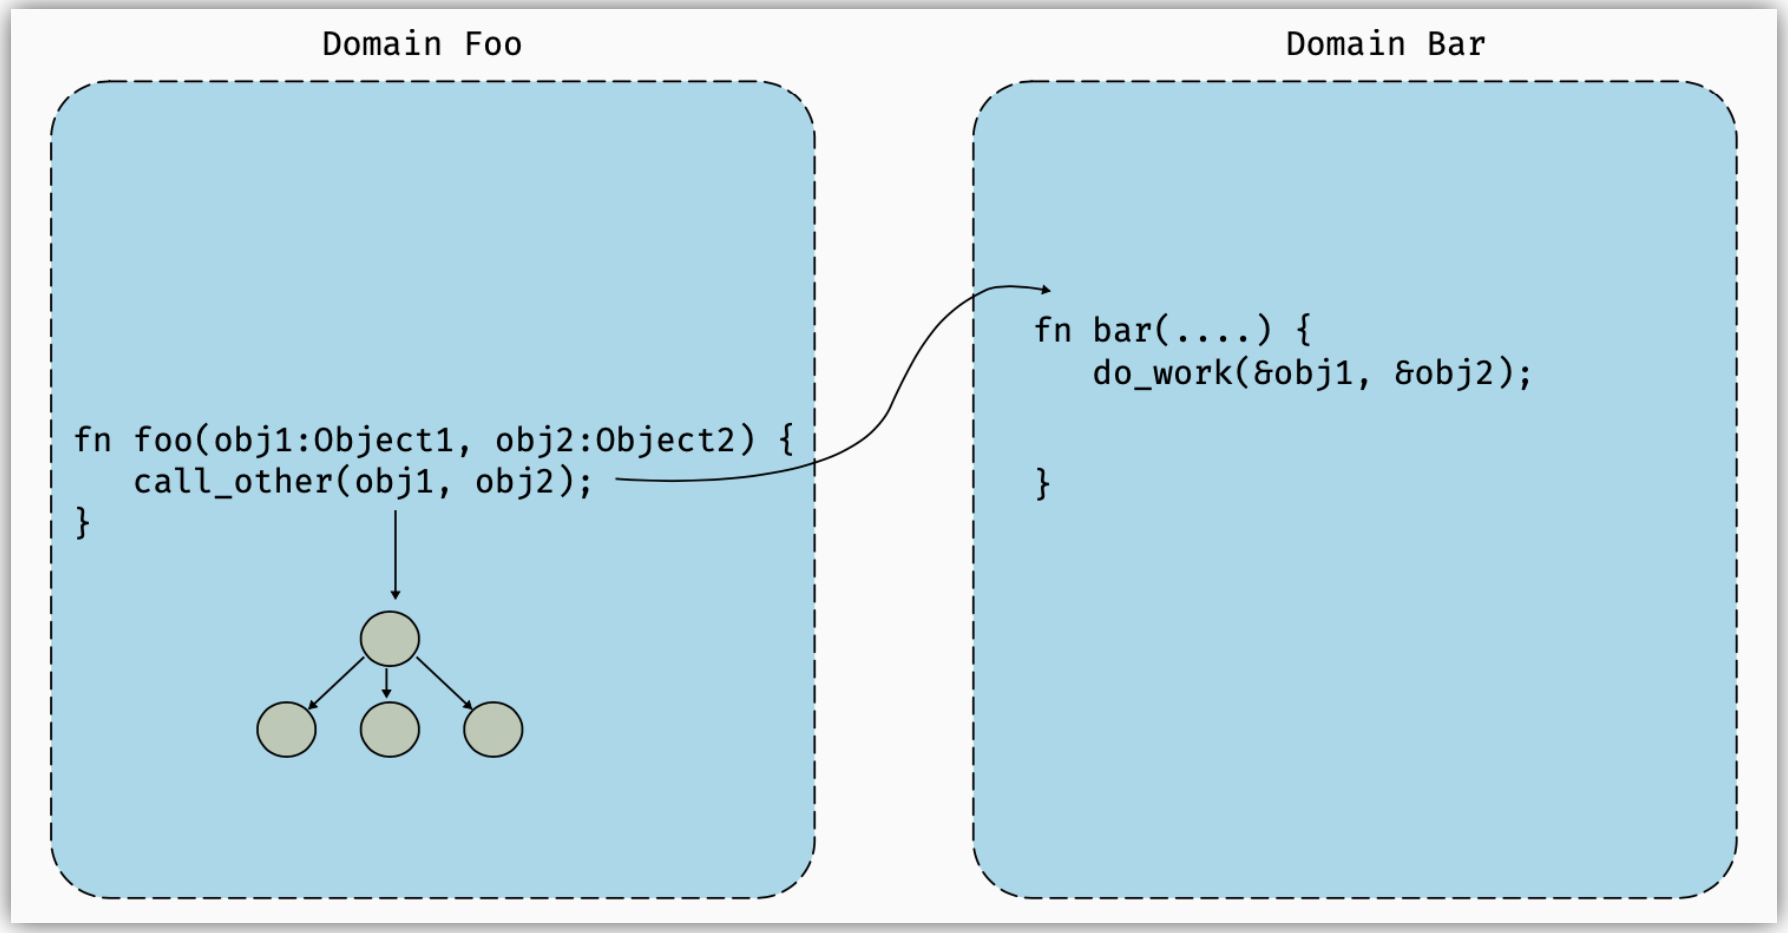

2.6 Fault Isolation in Language-based Systems: Deep Copy

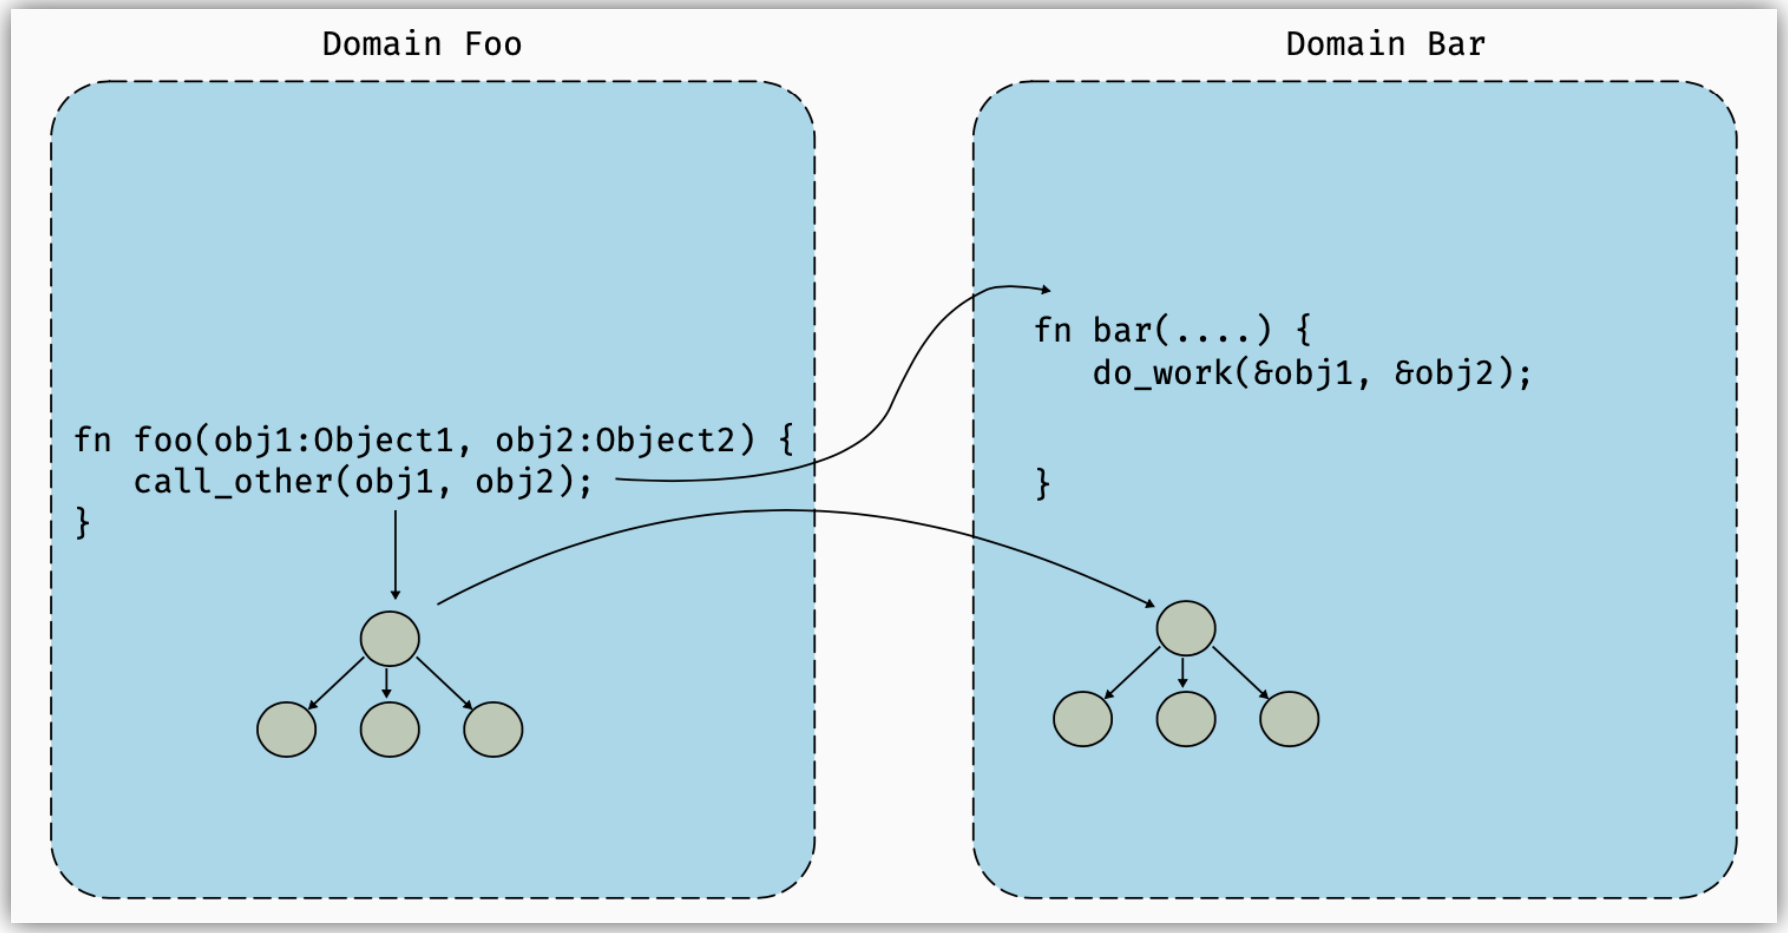

2.6 Fault Isolation in Language-based Systems: Deep Copy

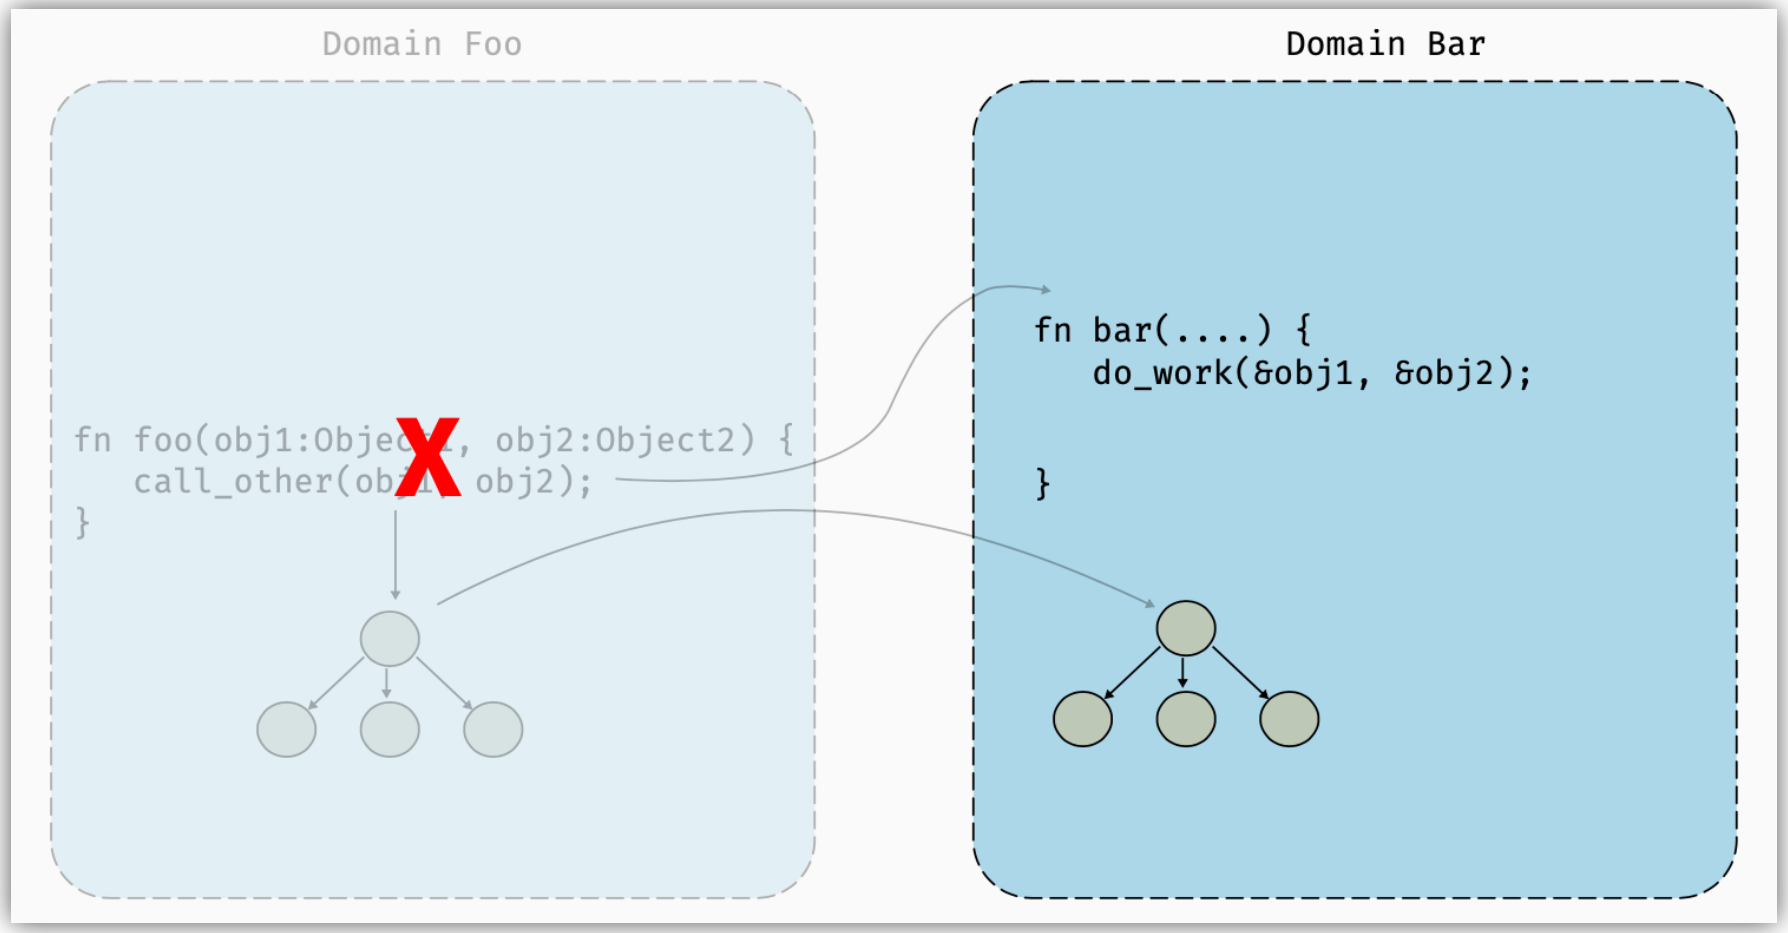

2.6 Fault Isolation in Language-based Systems: Deep Copy

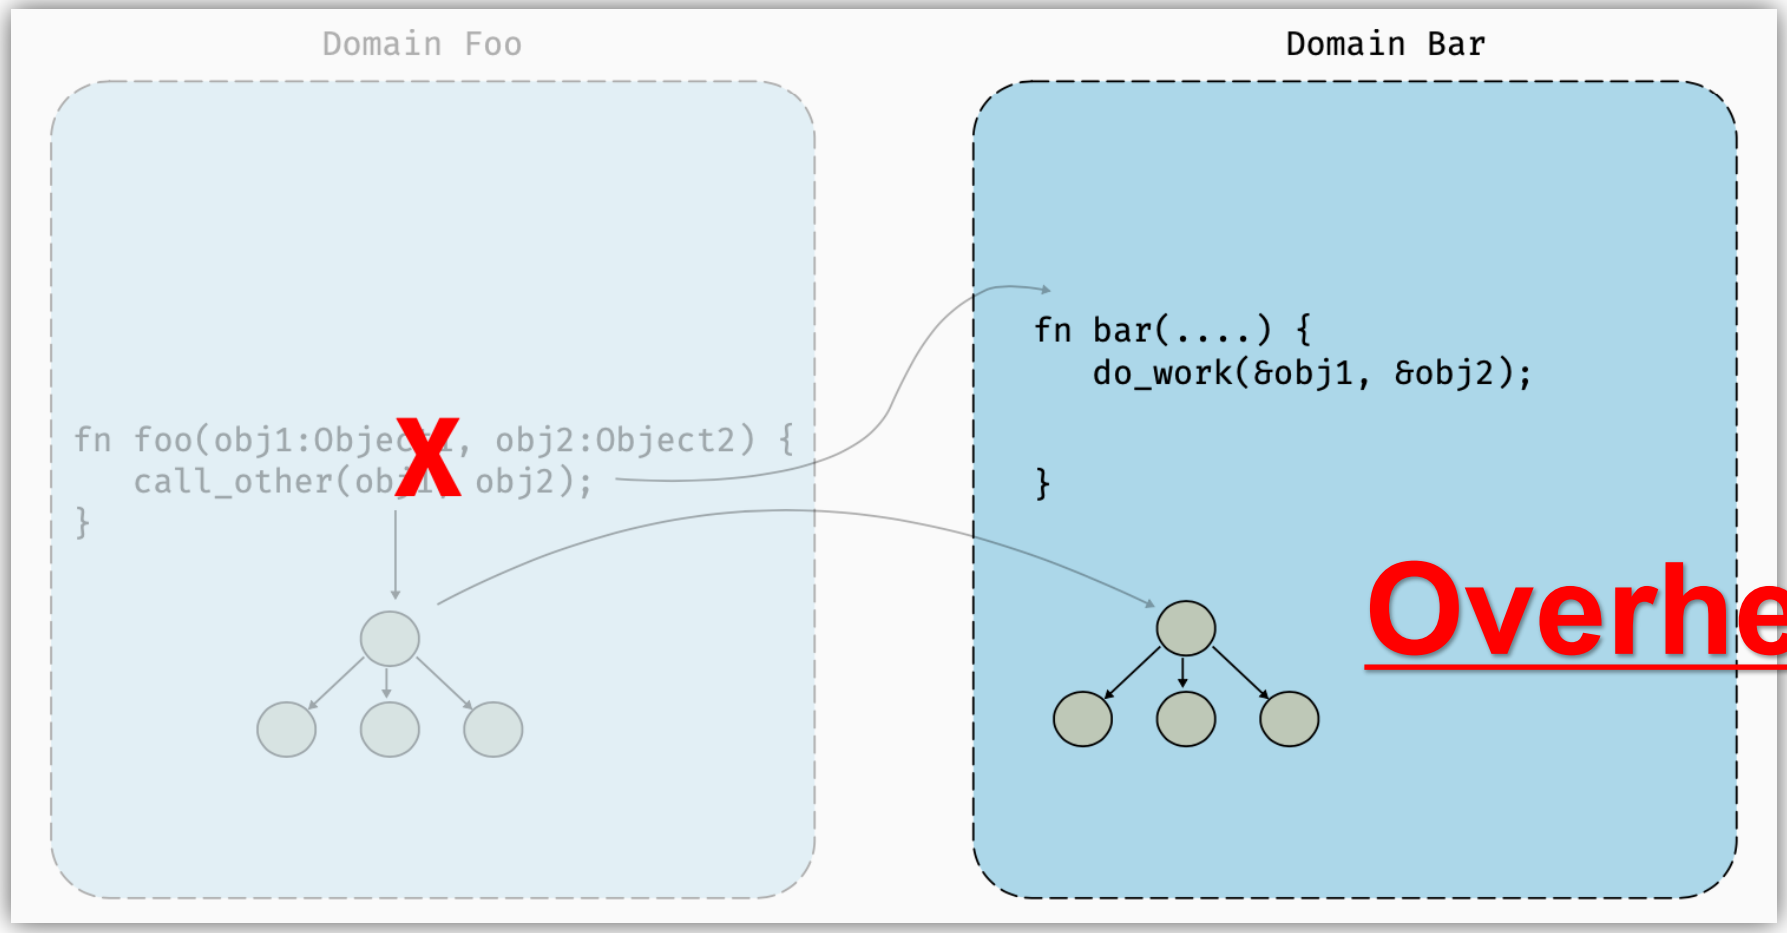

2.6 Fault Isolation in Language-based Systems: Singularity

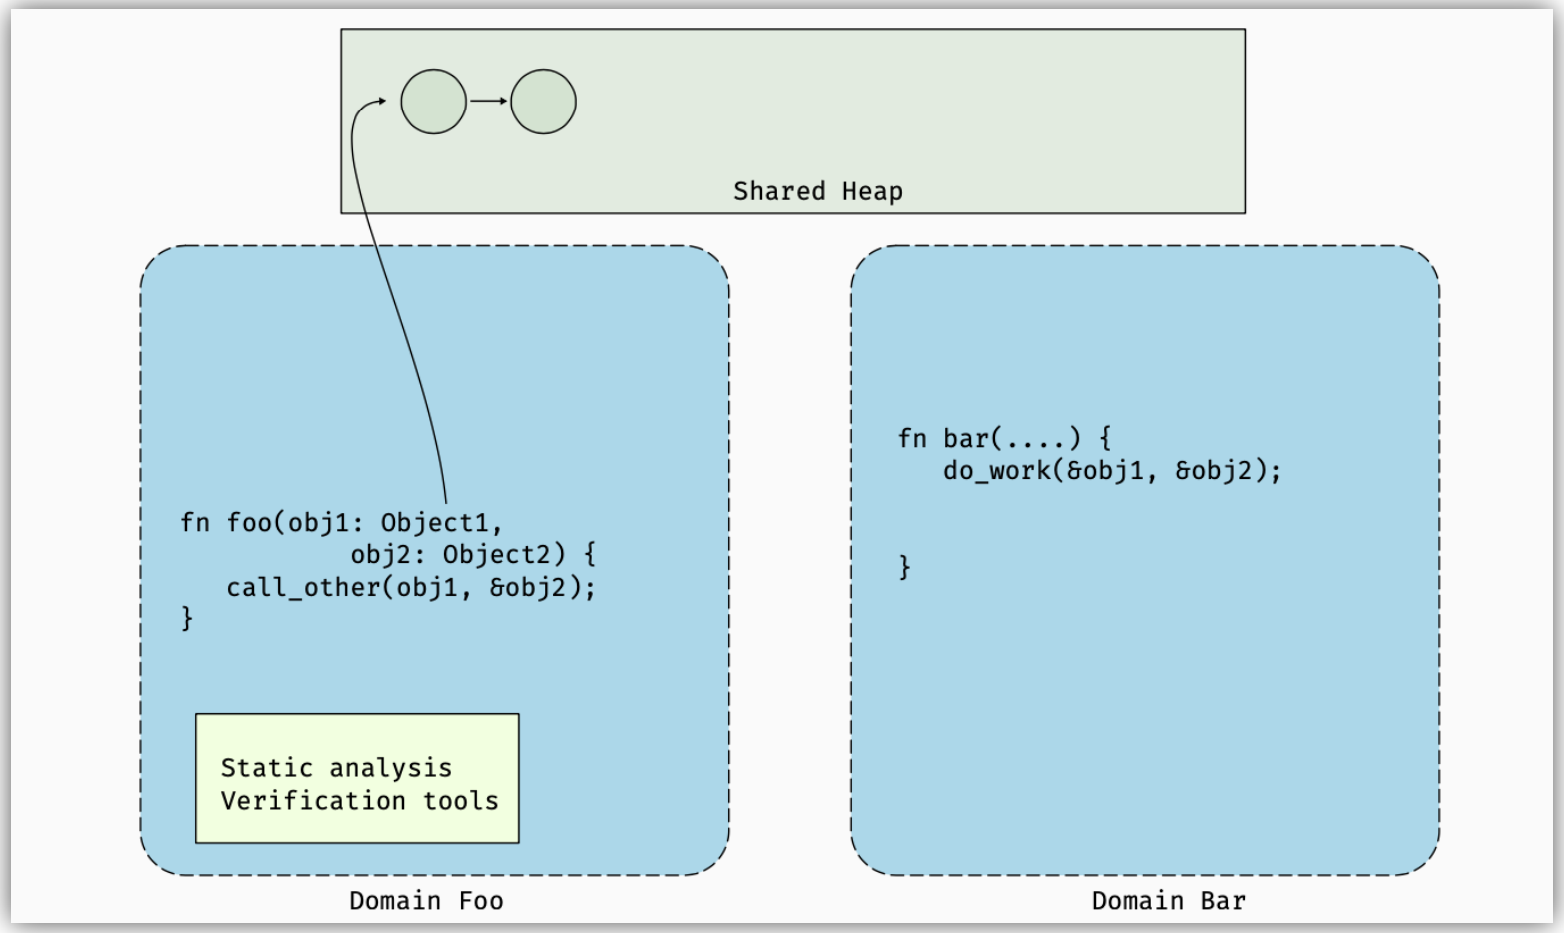

2.6 Fault Isolation in Language-based Systems: Singularity

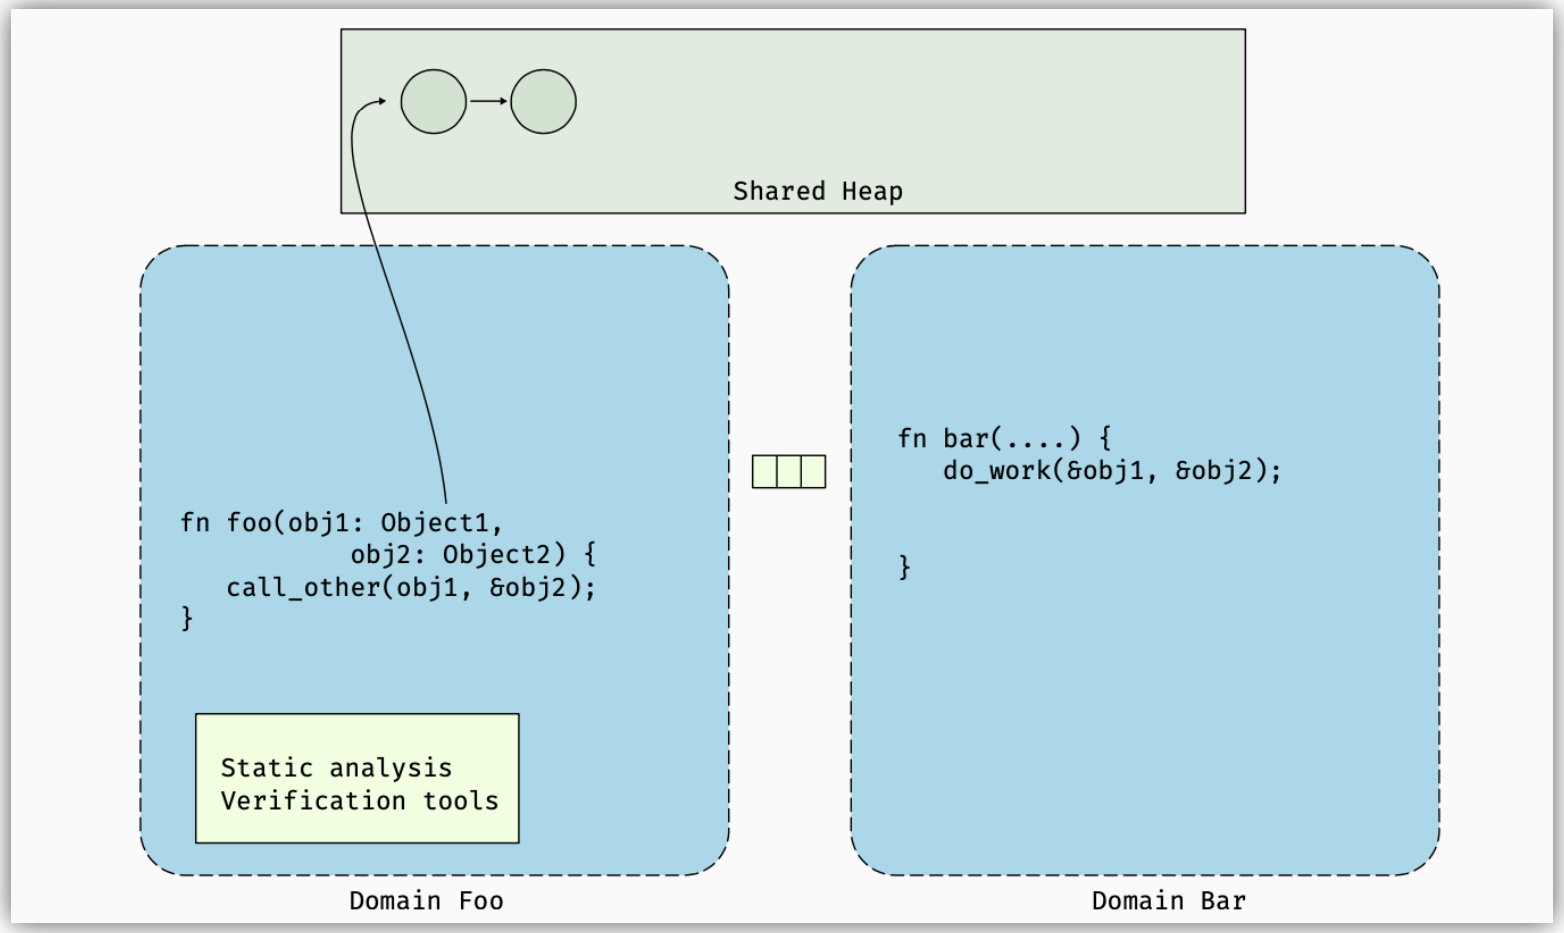

2.6 Fault Isolation in Language-based Systems: Singularity

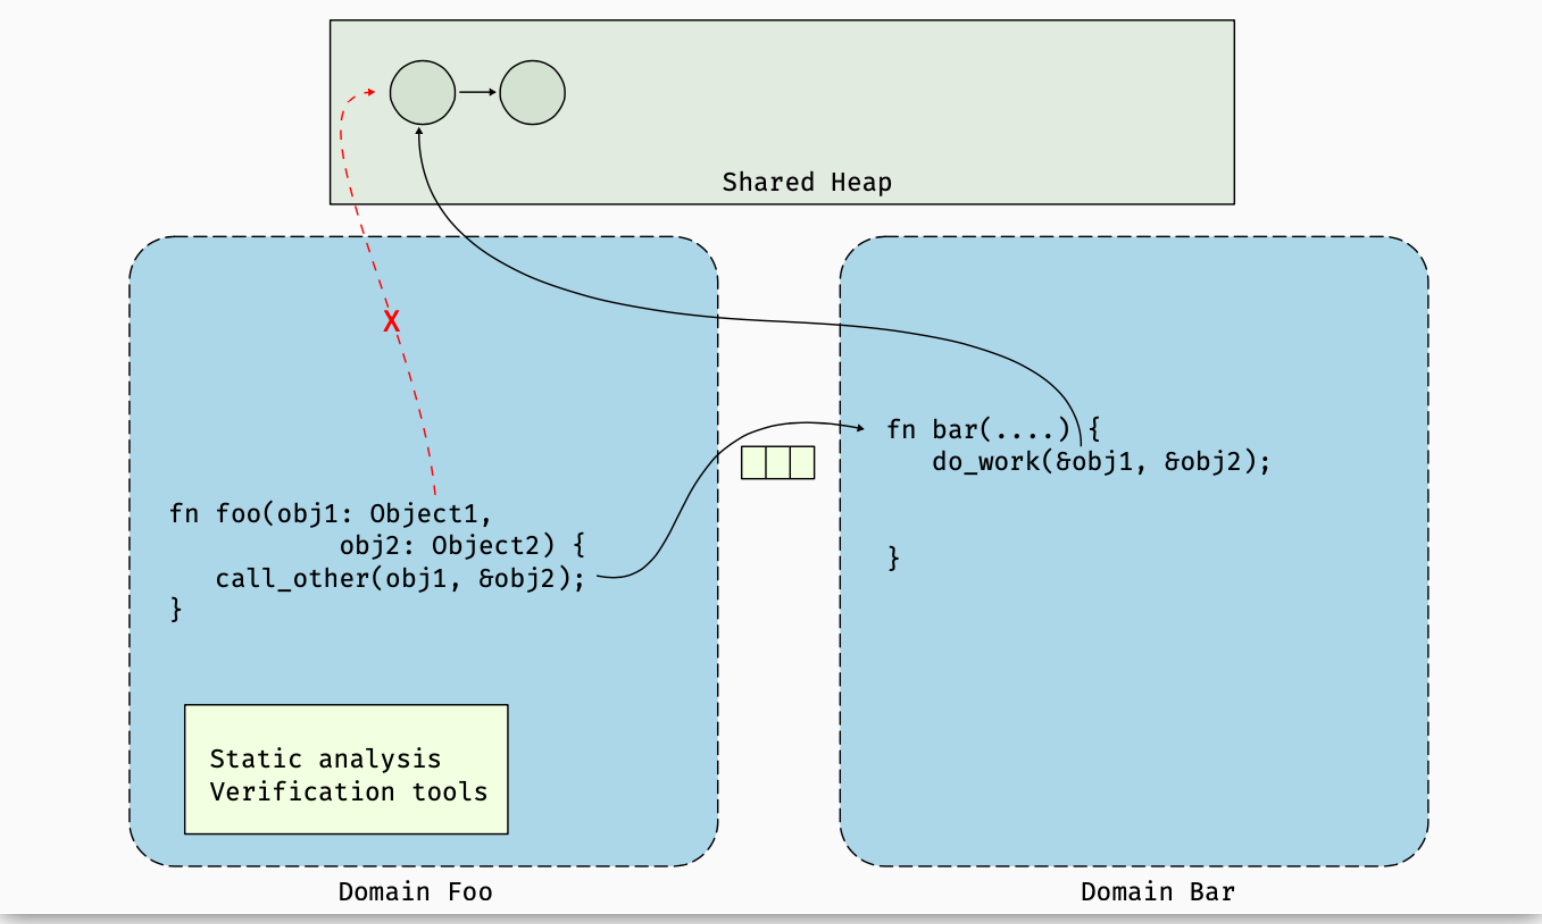

2.6 Fault Isolation in Language-based Systems: Singularity

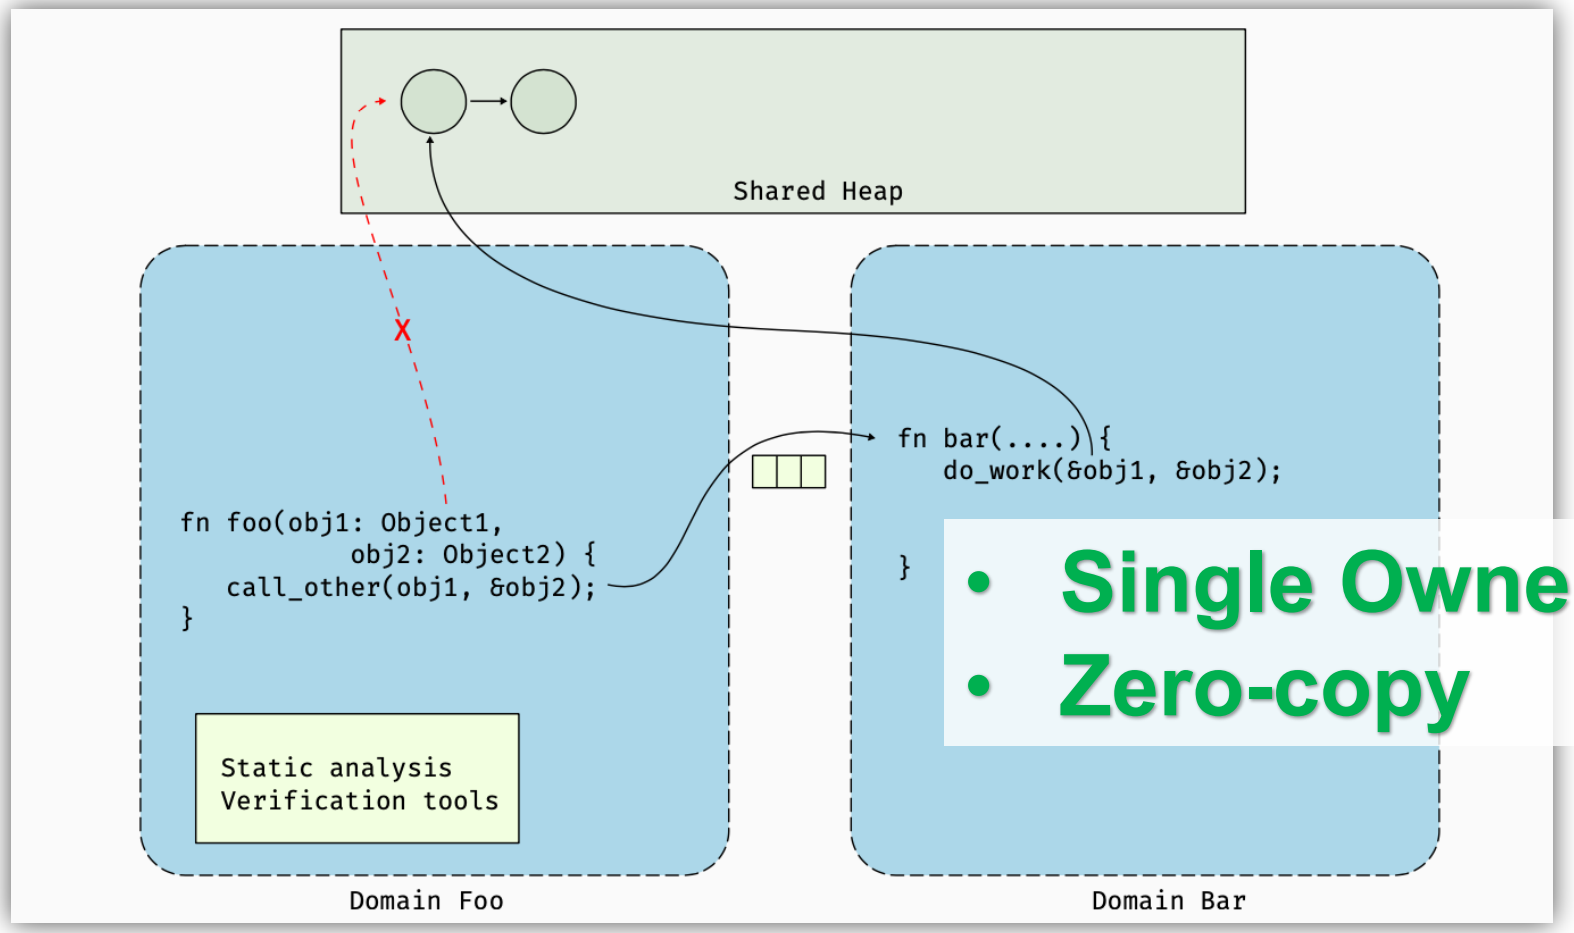

- Single Ownership
- Zero-copy

2.7 Summary of Background

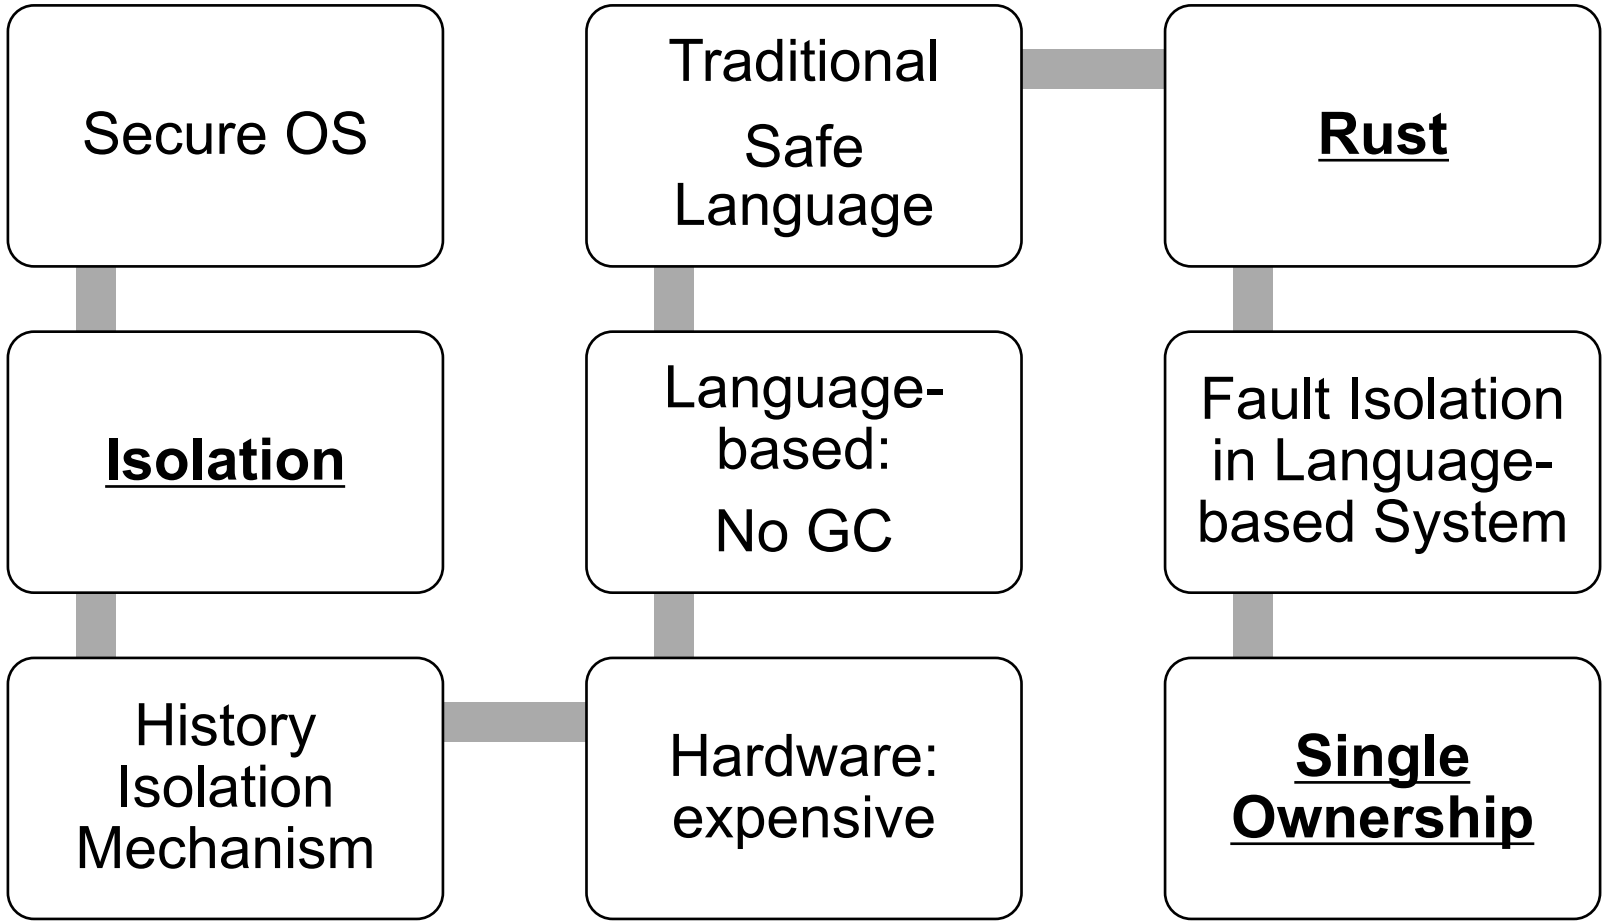

3. RedLeaf

3.1 Architecture

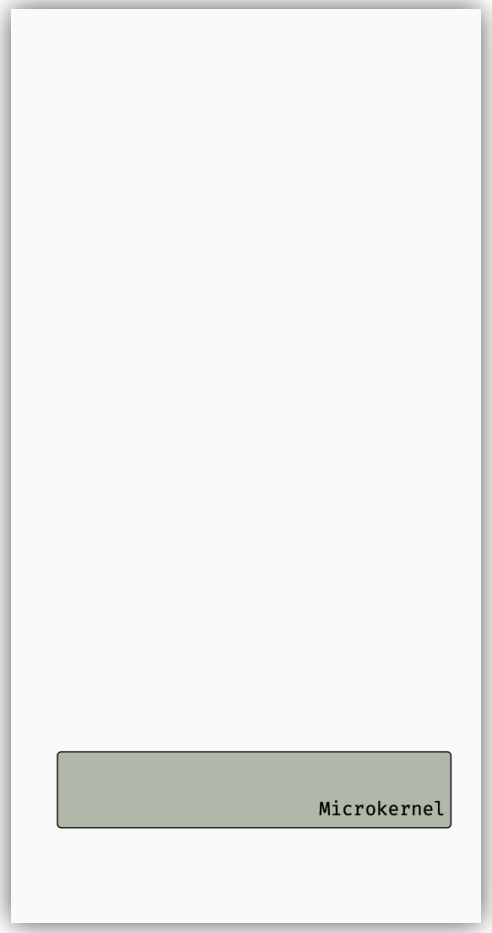

3.1 Architecture

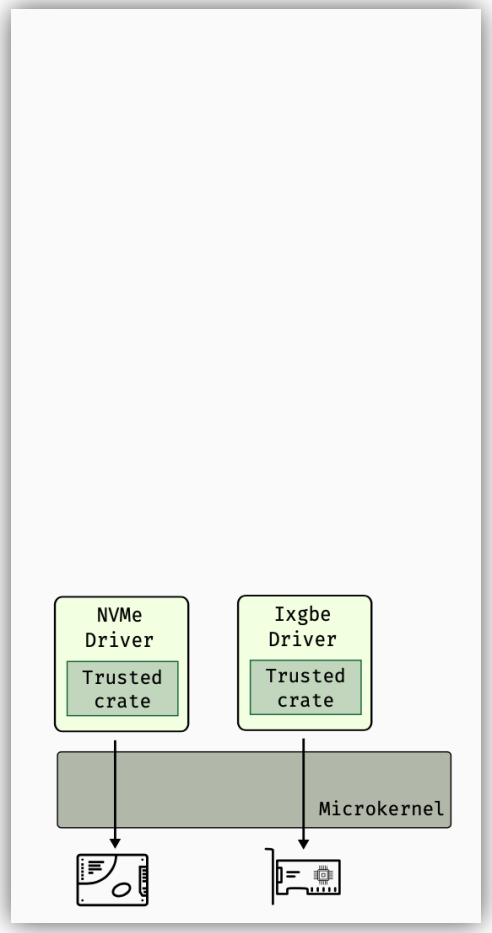

3.1 Architecture

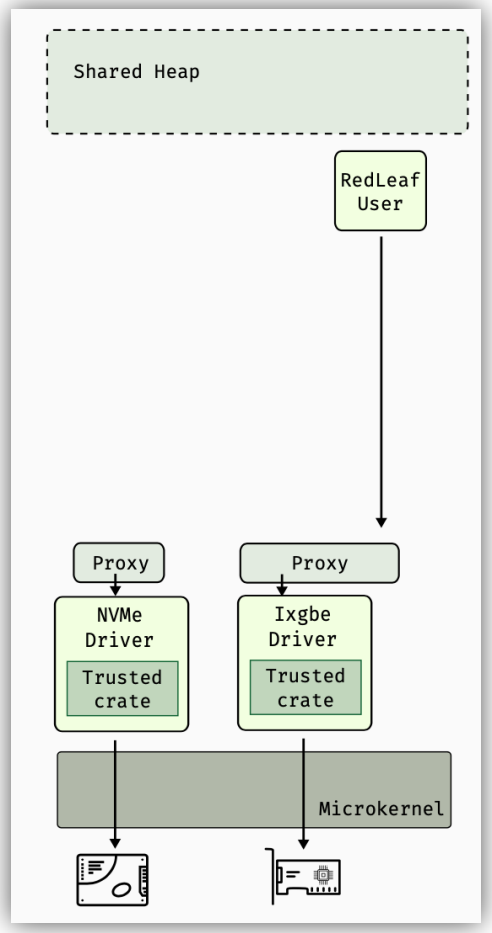

3.1 Architecture

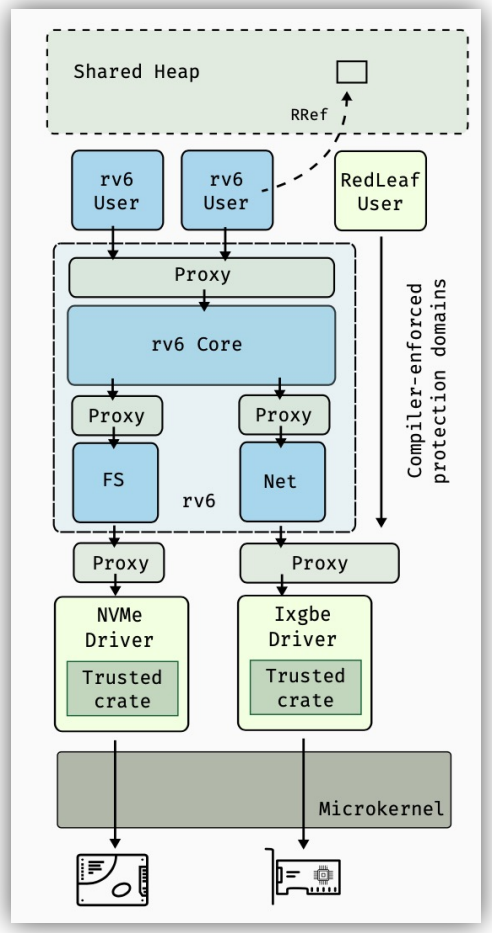

3.2 * Trust Base

- Rust **compiler**
- Rust core **libraries** (crates)
- Non-malicious devices (can be spared by IOMMU)

3.3 Fault Isolation

- After a domain crash
  - Unwind all threads running inside
  - Subsequent invocations return error
  - All resources are deallocated
  - Other threads continue execution

3.4 Heap Isolation

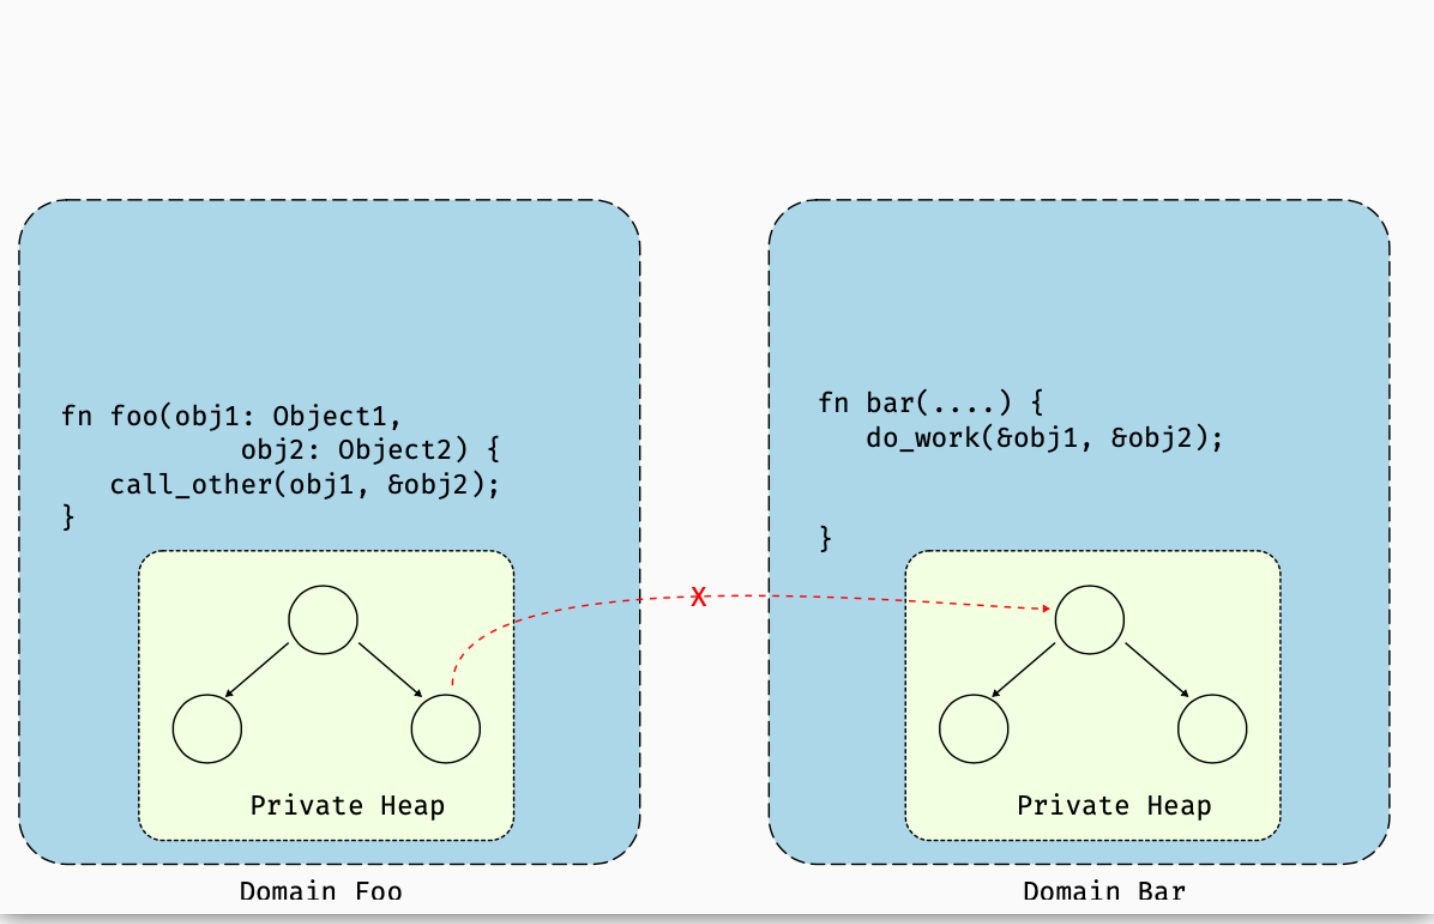

3.4 Heap Isolation

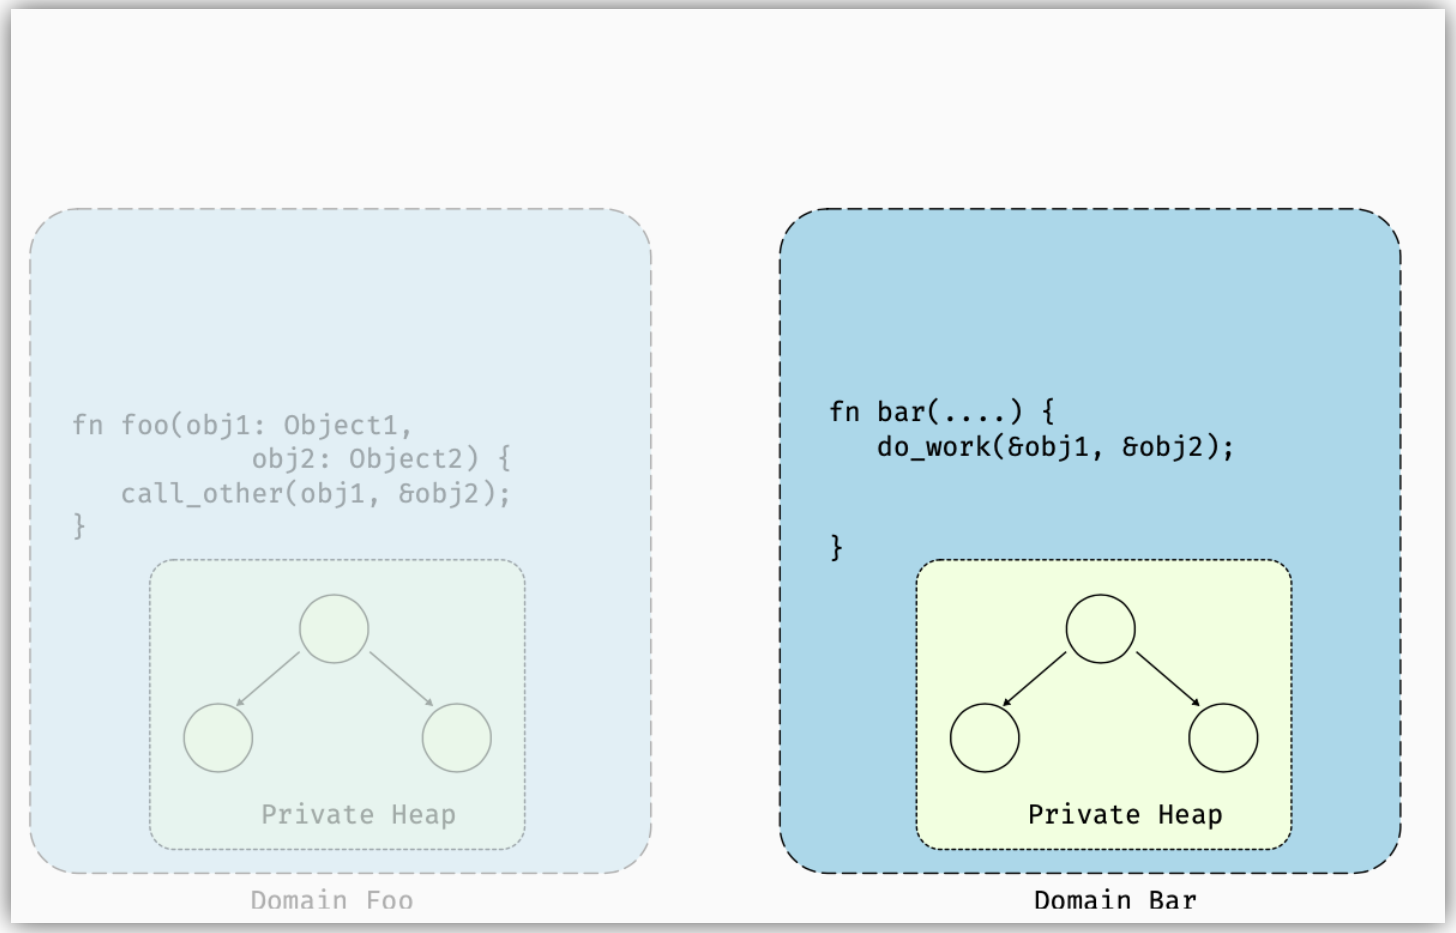

3.4 Heap Isolation

Domains never hold pointers to other domains

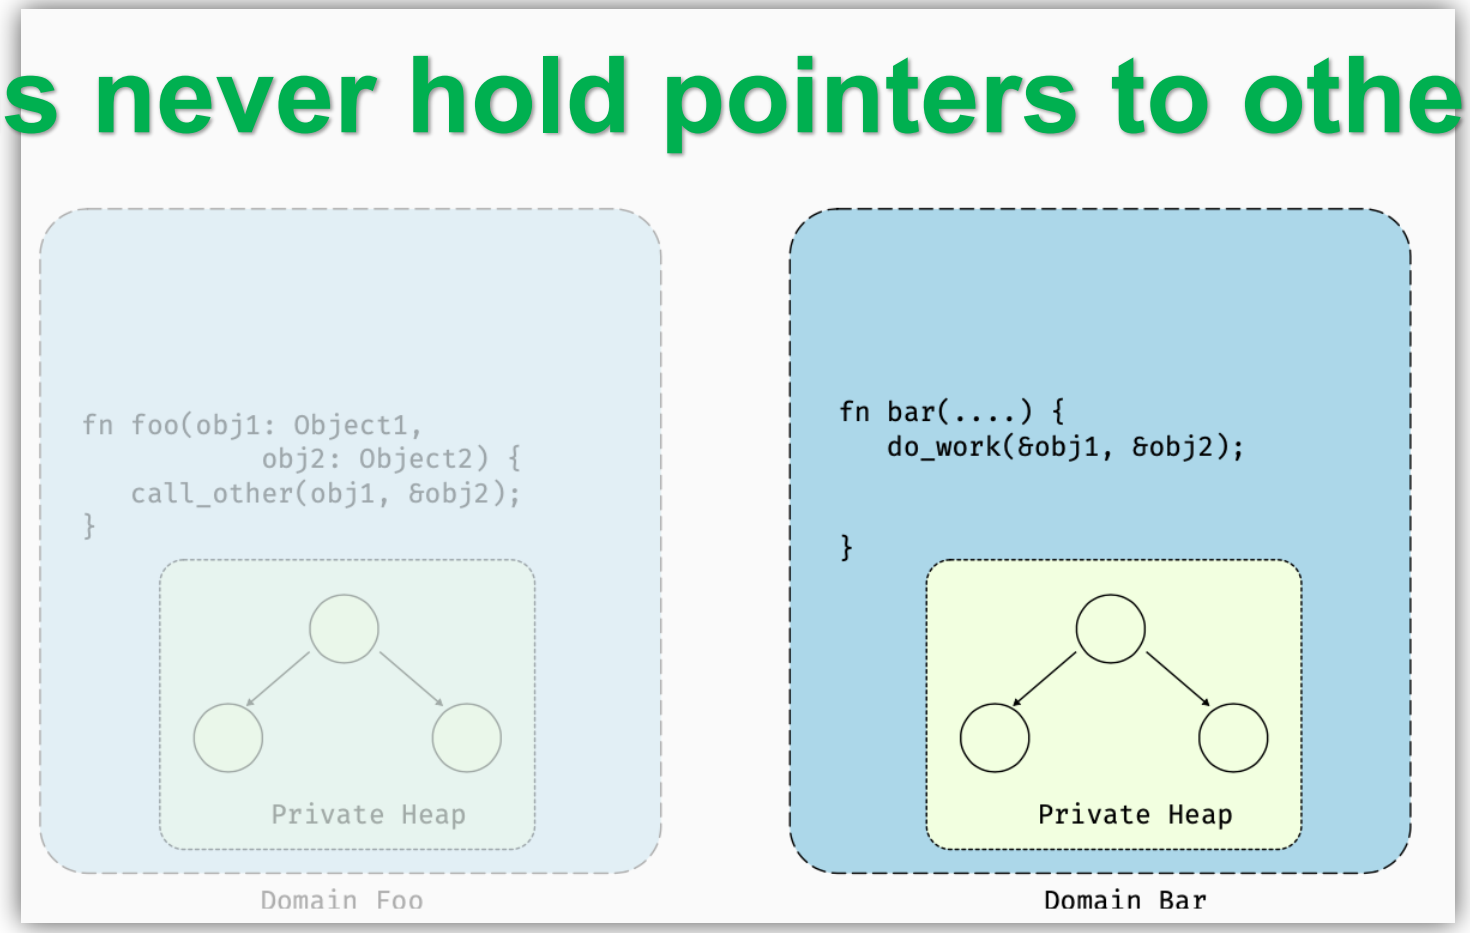

3.4 Heap Isolation

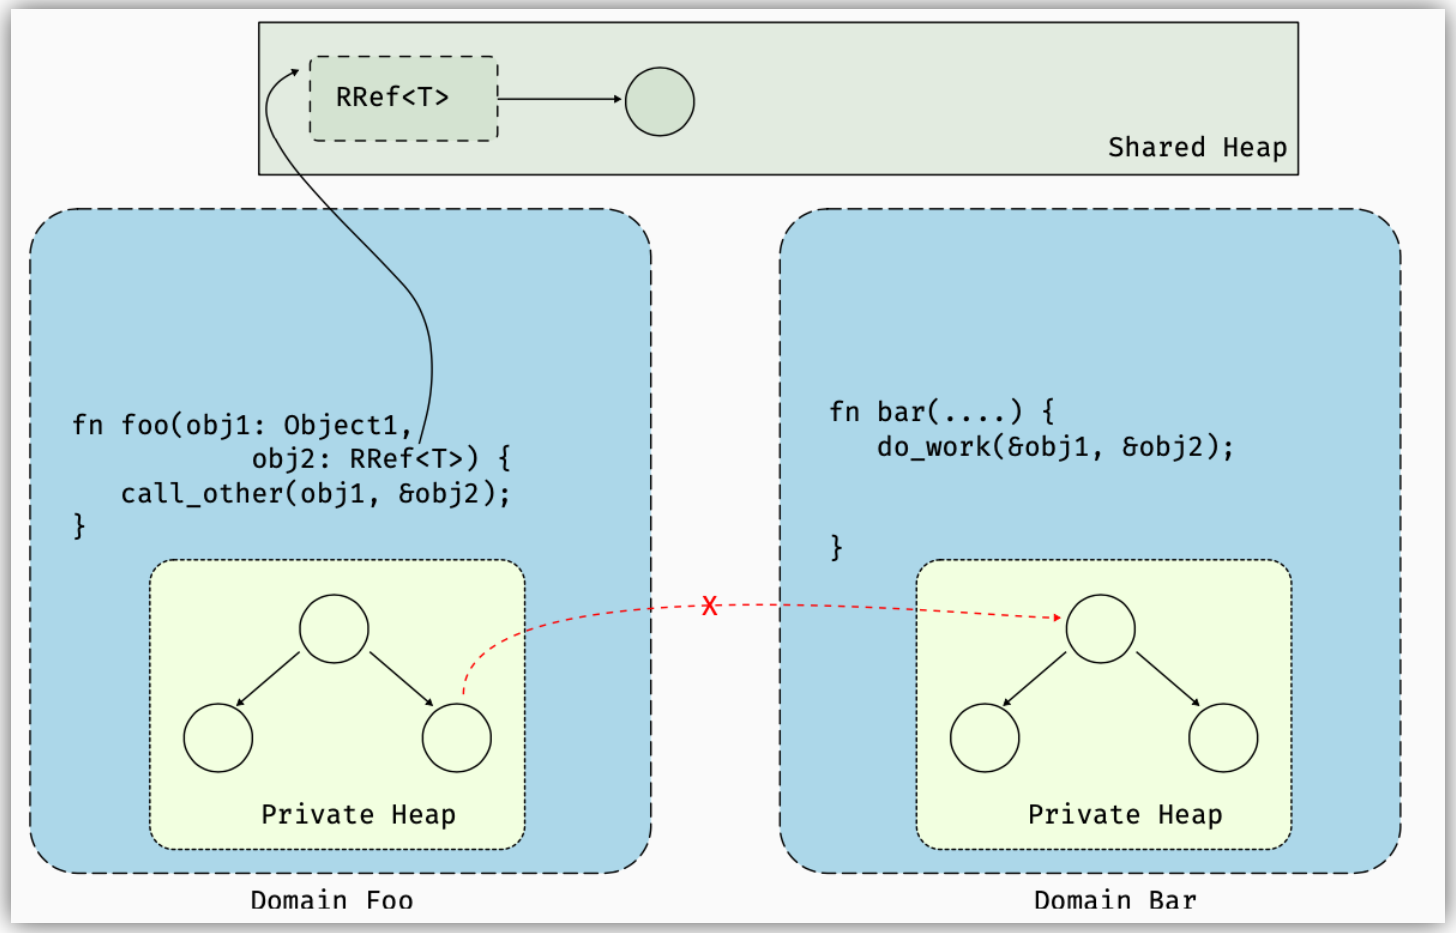

3.4 Heap Isolation

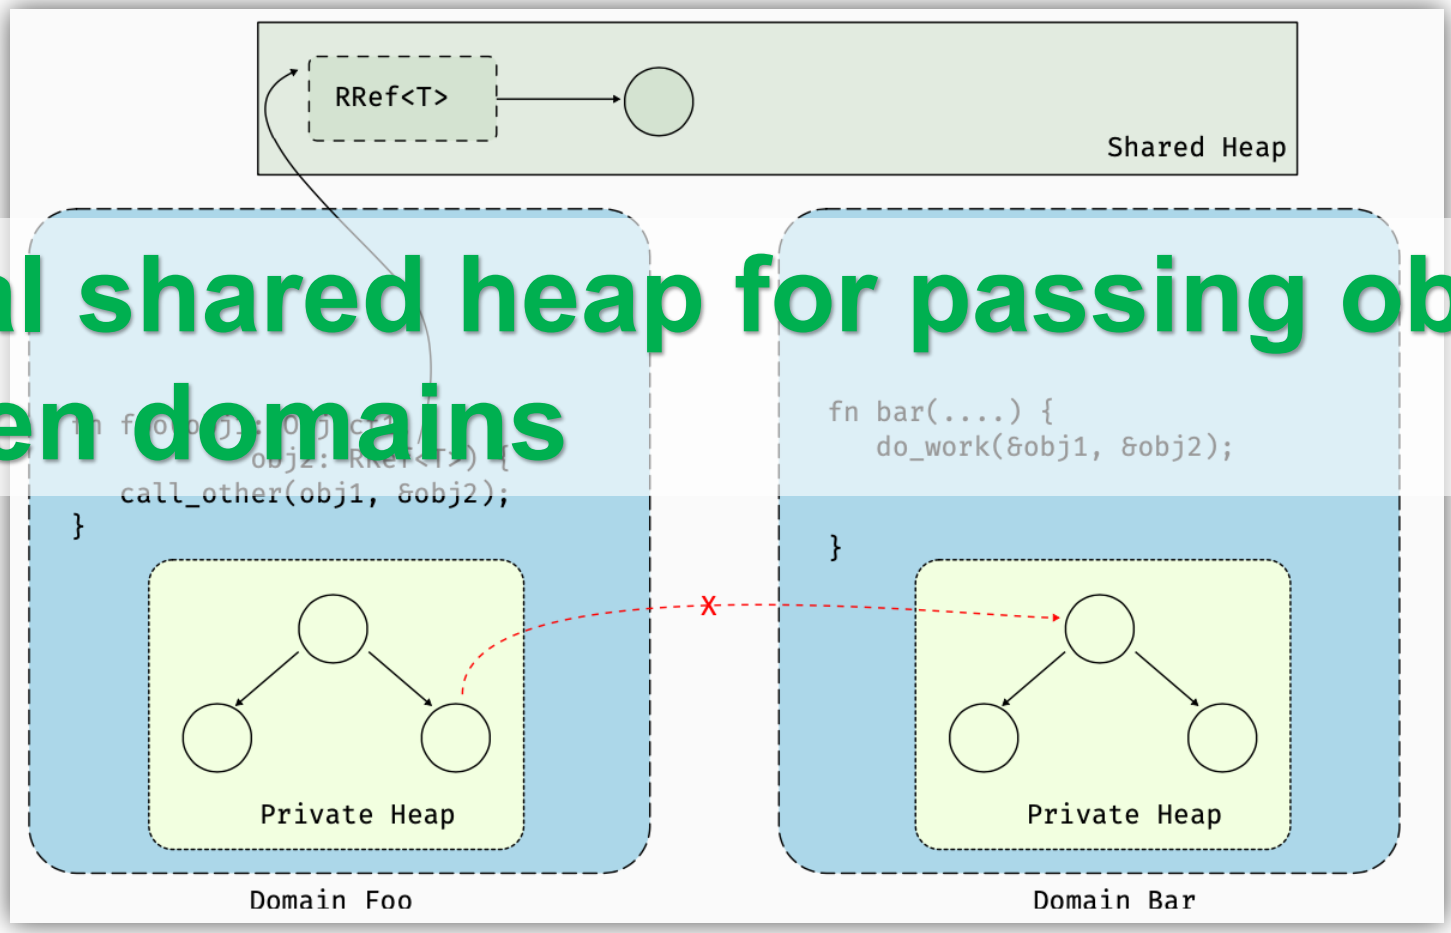

Special shared heap for passing objects between domains

3.4 Heap Isolation

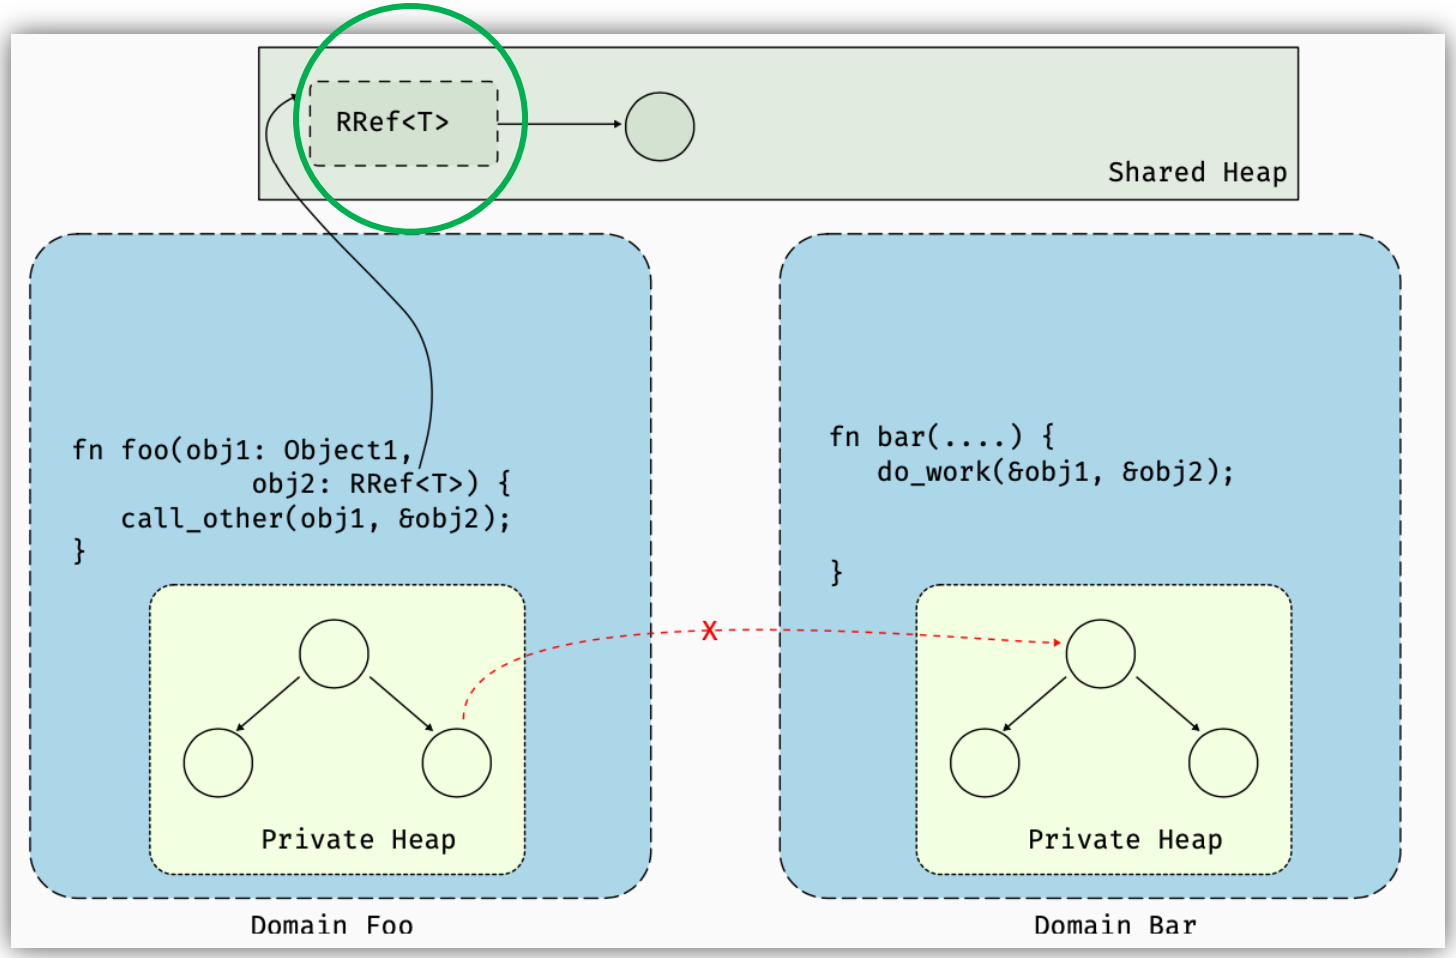

3.4 Heap Isolation

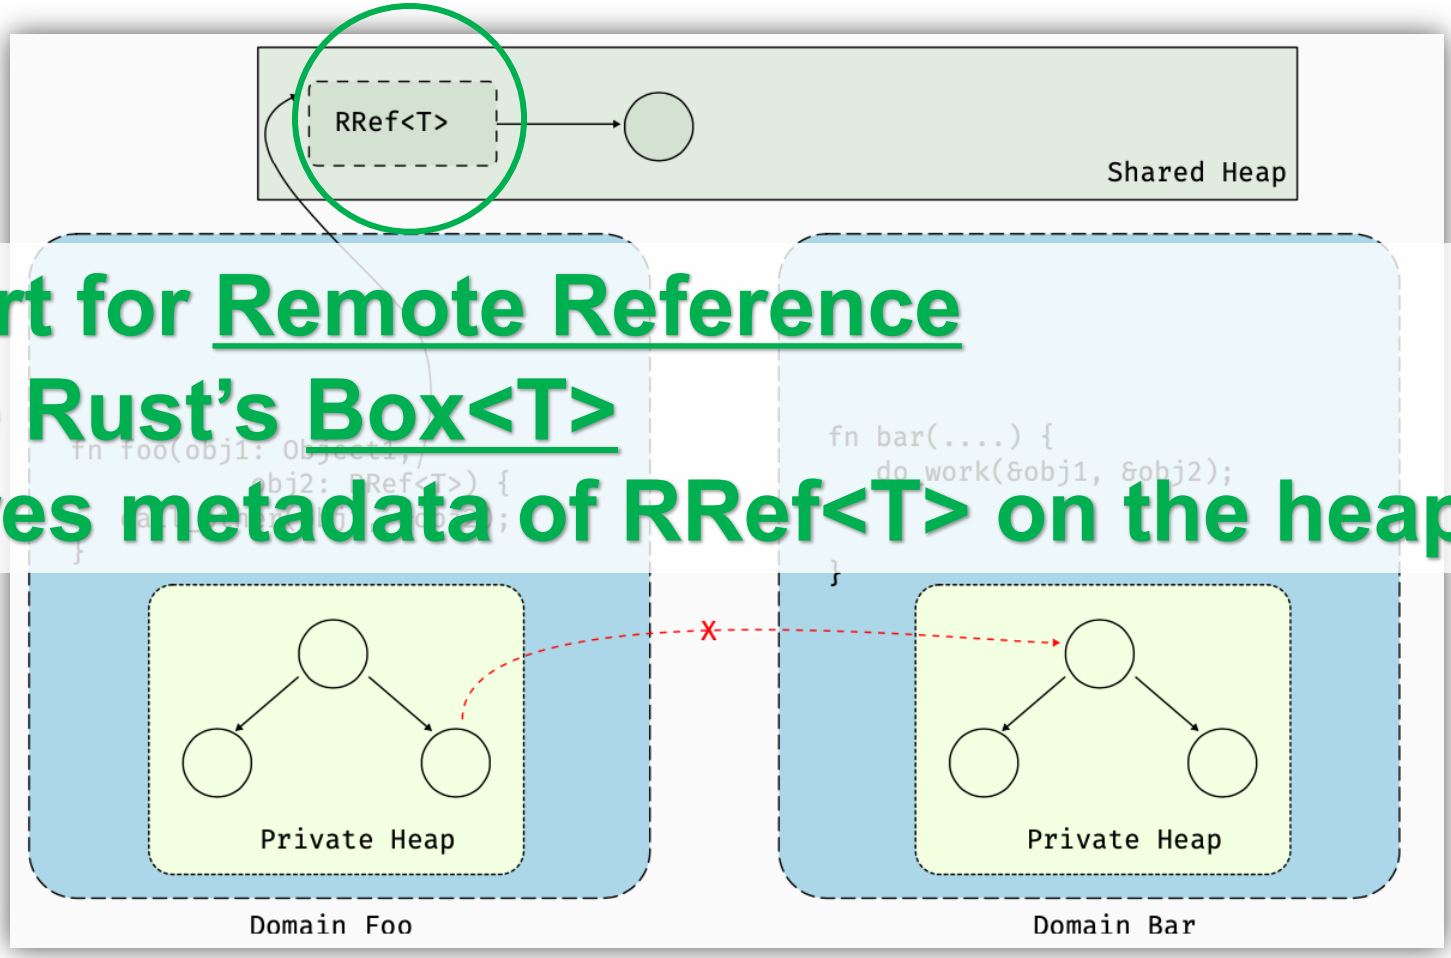

1. Short for Remote Reference
2. Like Rust's Box<T>
3. Stores metadata of RRef<T> on the heap

3.5 Exchangeable Types

- Exchangeable Types can be:
  - **RRef<T>** itself
  - A subset of Rust primitive **Copy** types (not references or pointers)
  - **Composite types** constructed out of exchangeable types
  - References to traits **with methods that receive exchangeable types**

3.5 Exchangeable Types

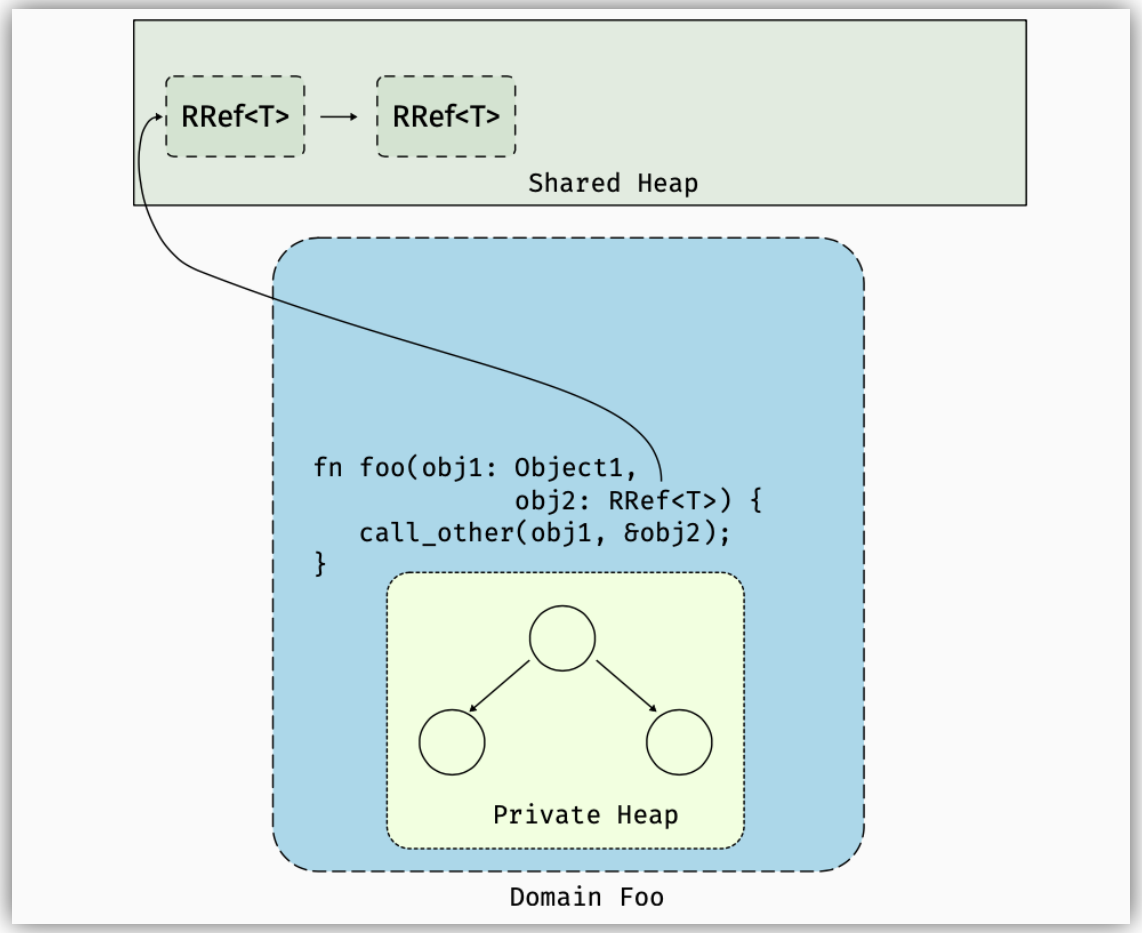

3.5 Exchangeable Types

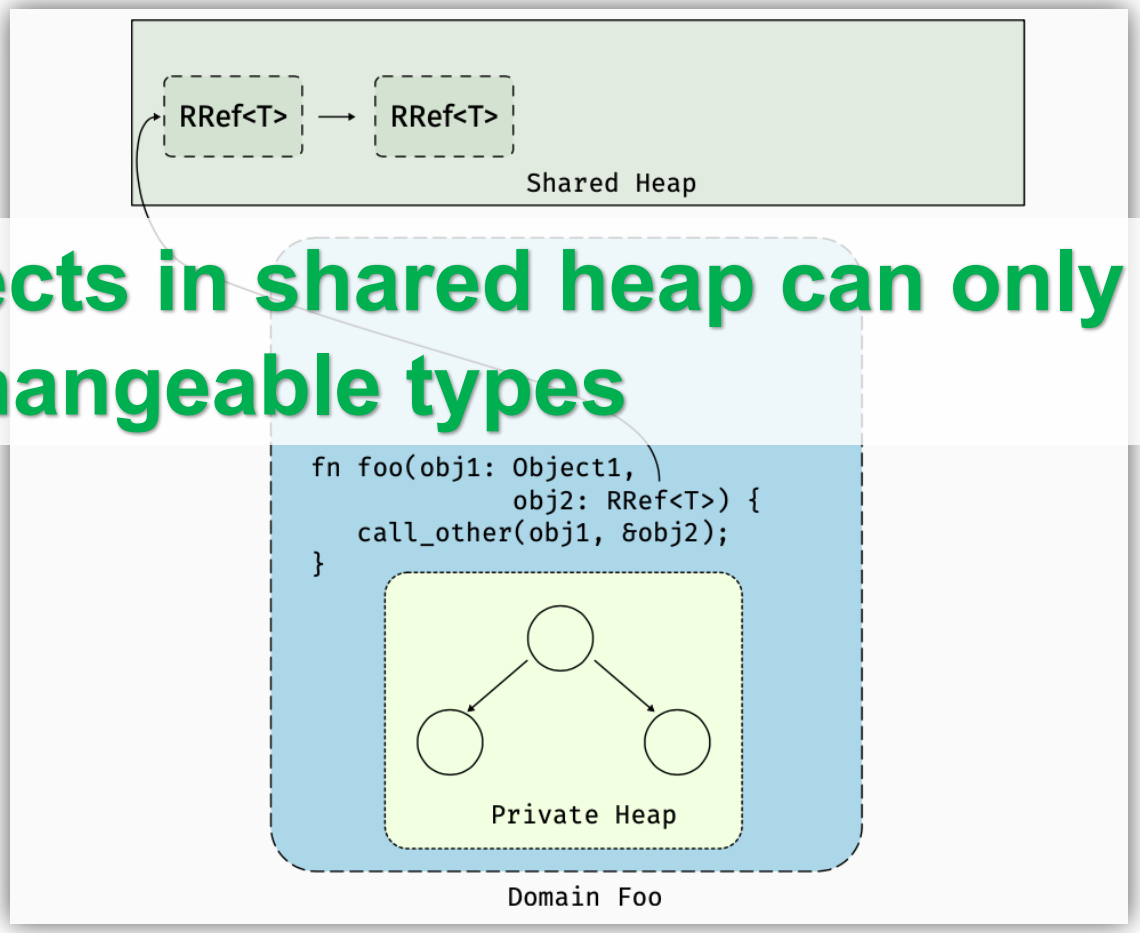

Objects in shared heap can only be exchangeable types

3.5 Exchangeable Types

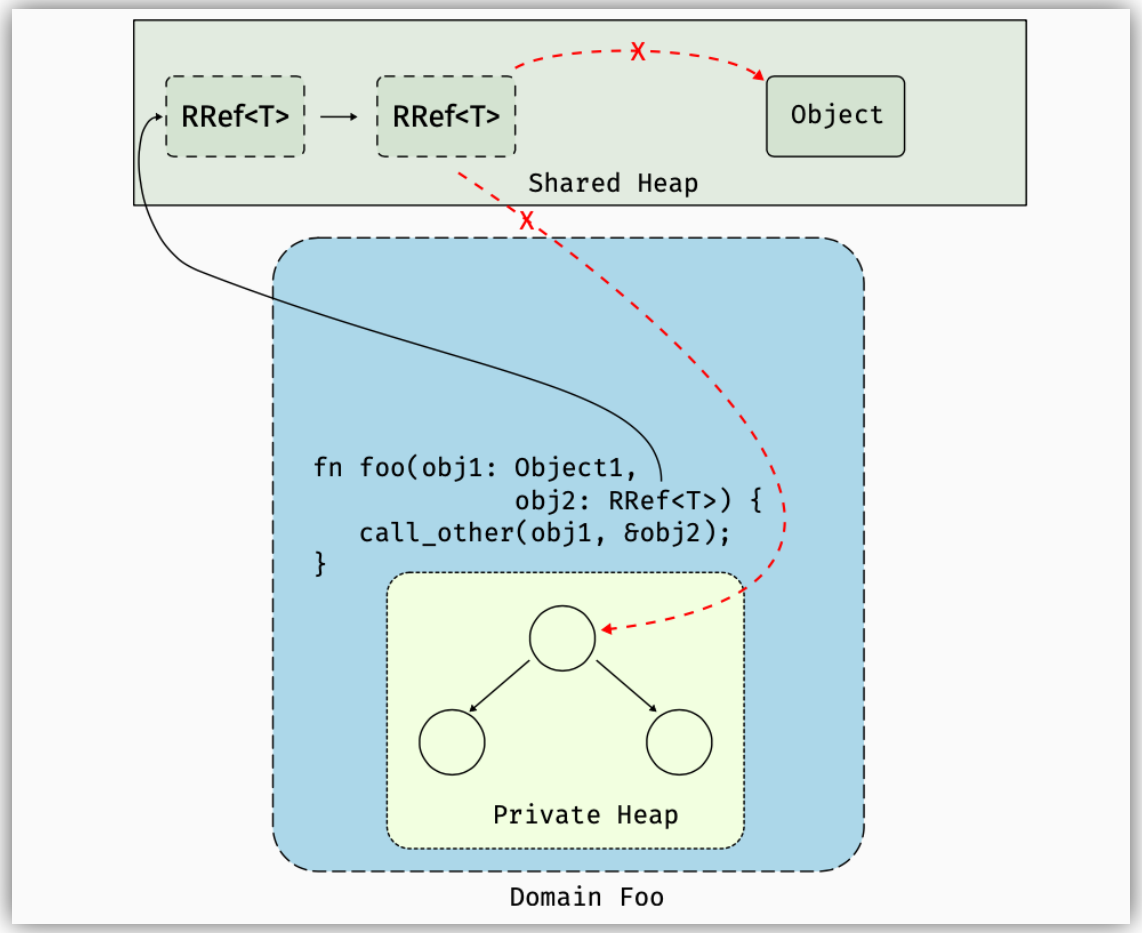

3.6 Ownership Tracking

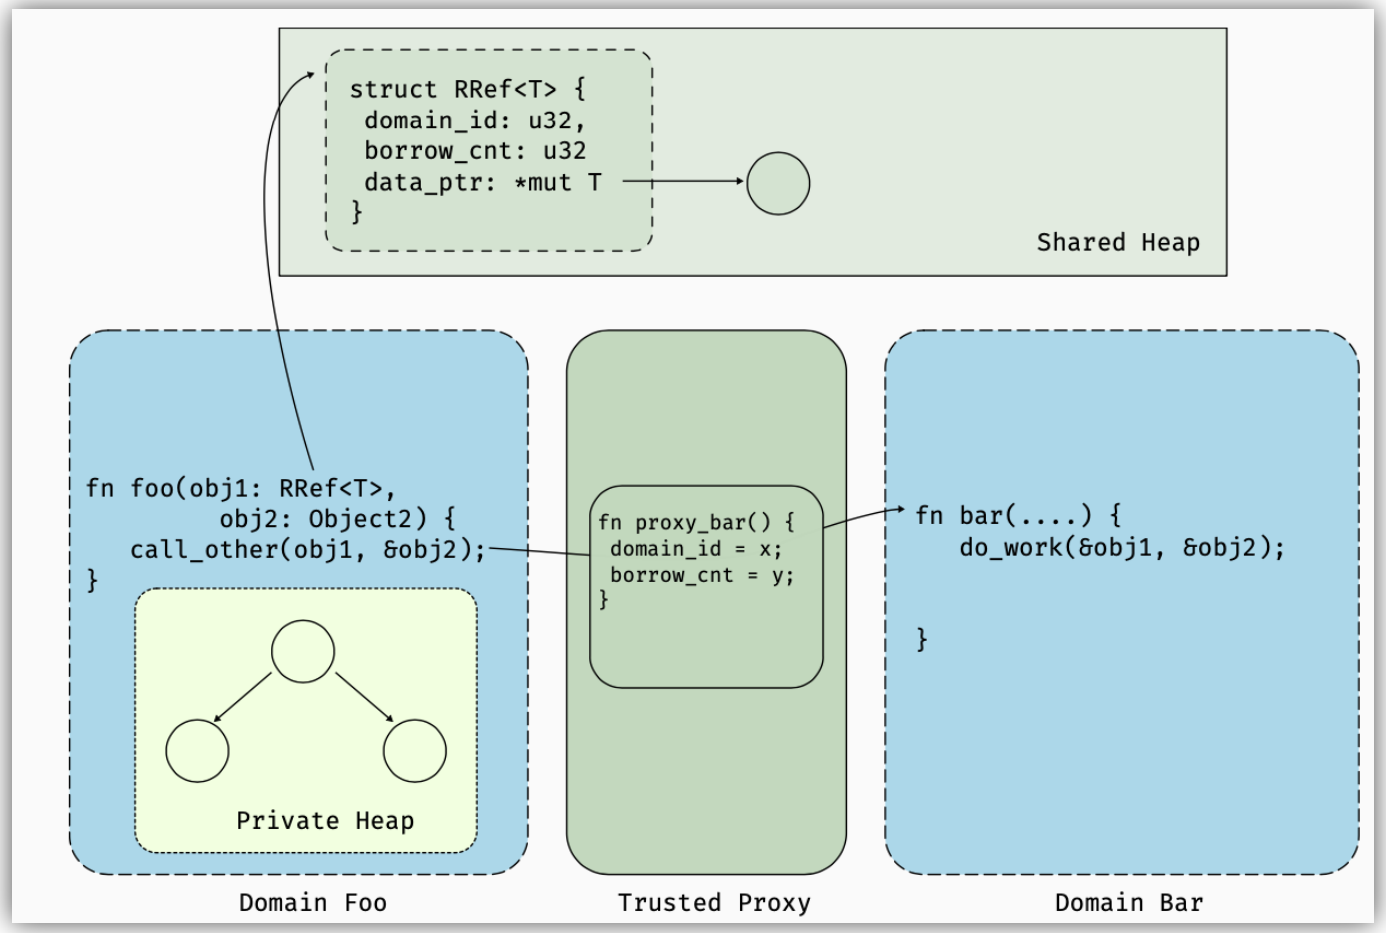

3.6 Ownership Tracking

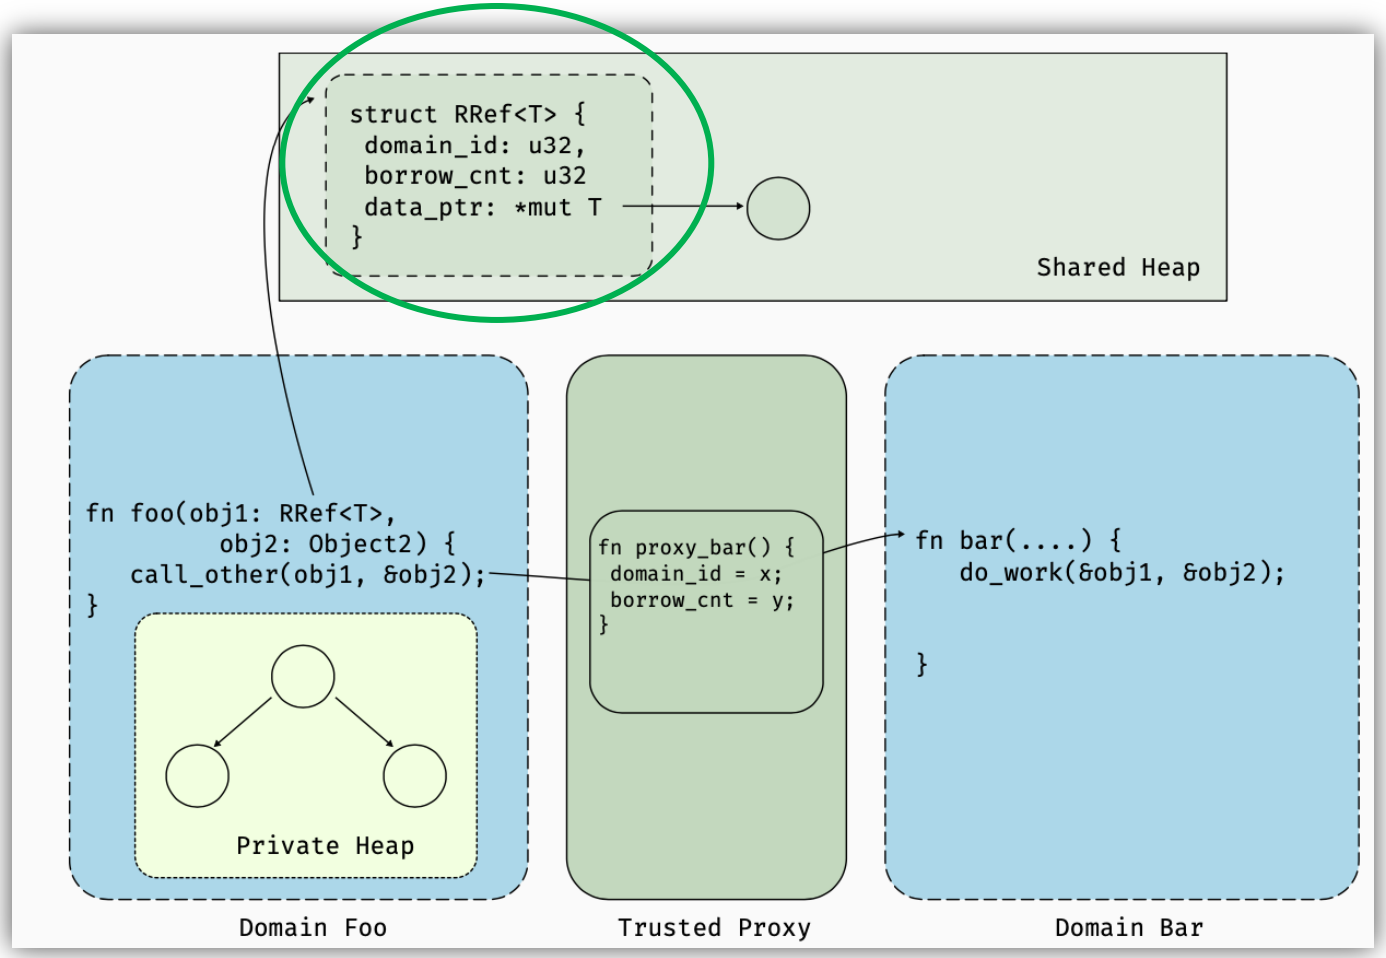

3.6 Ownership Tracking

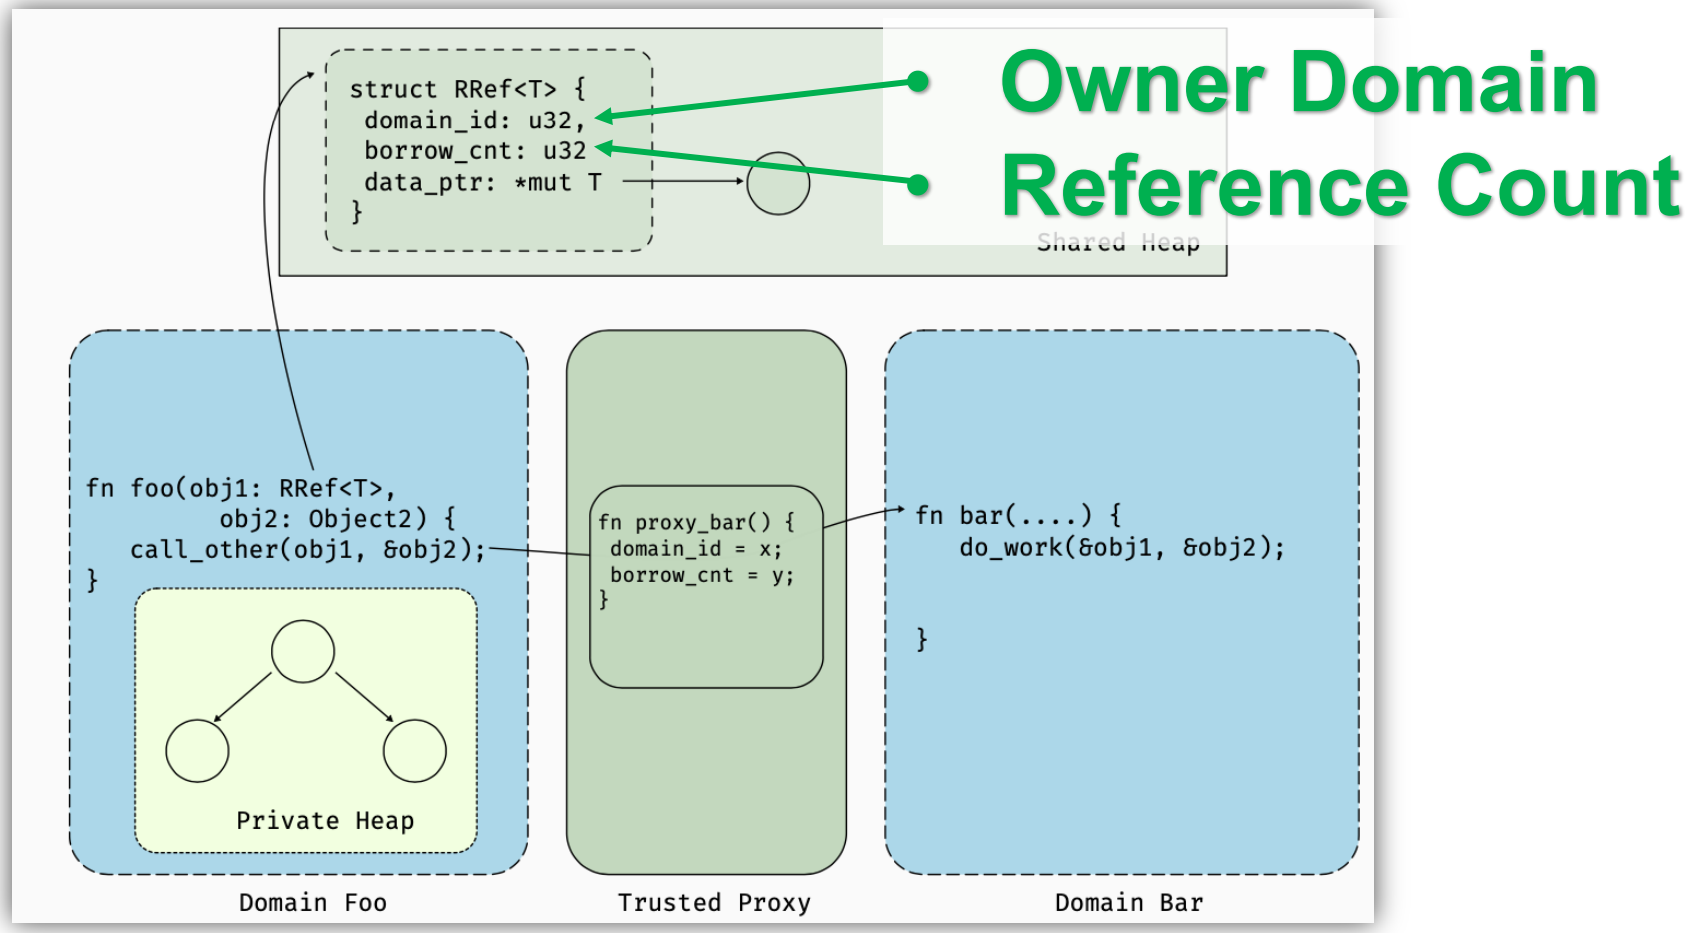

3.6 Ownership Tracking

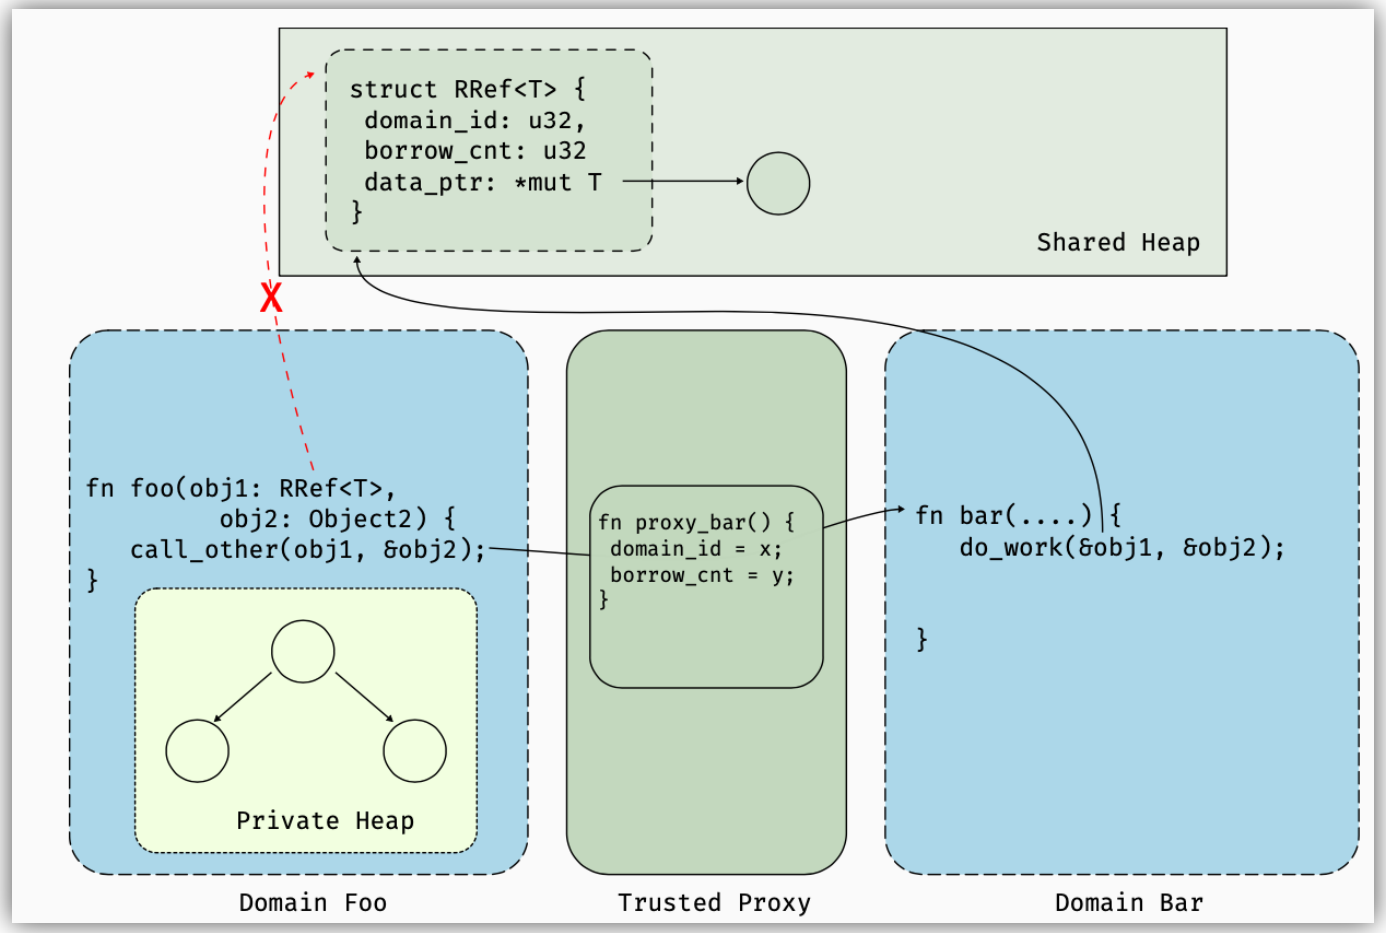

3.6 Ownership Tracking

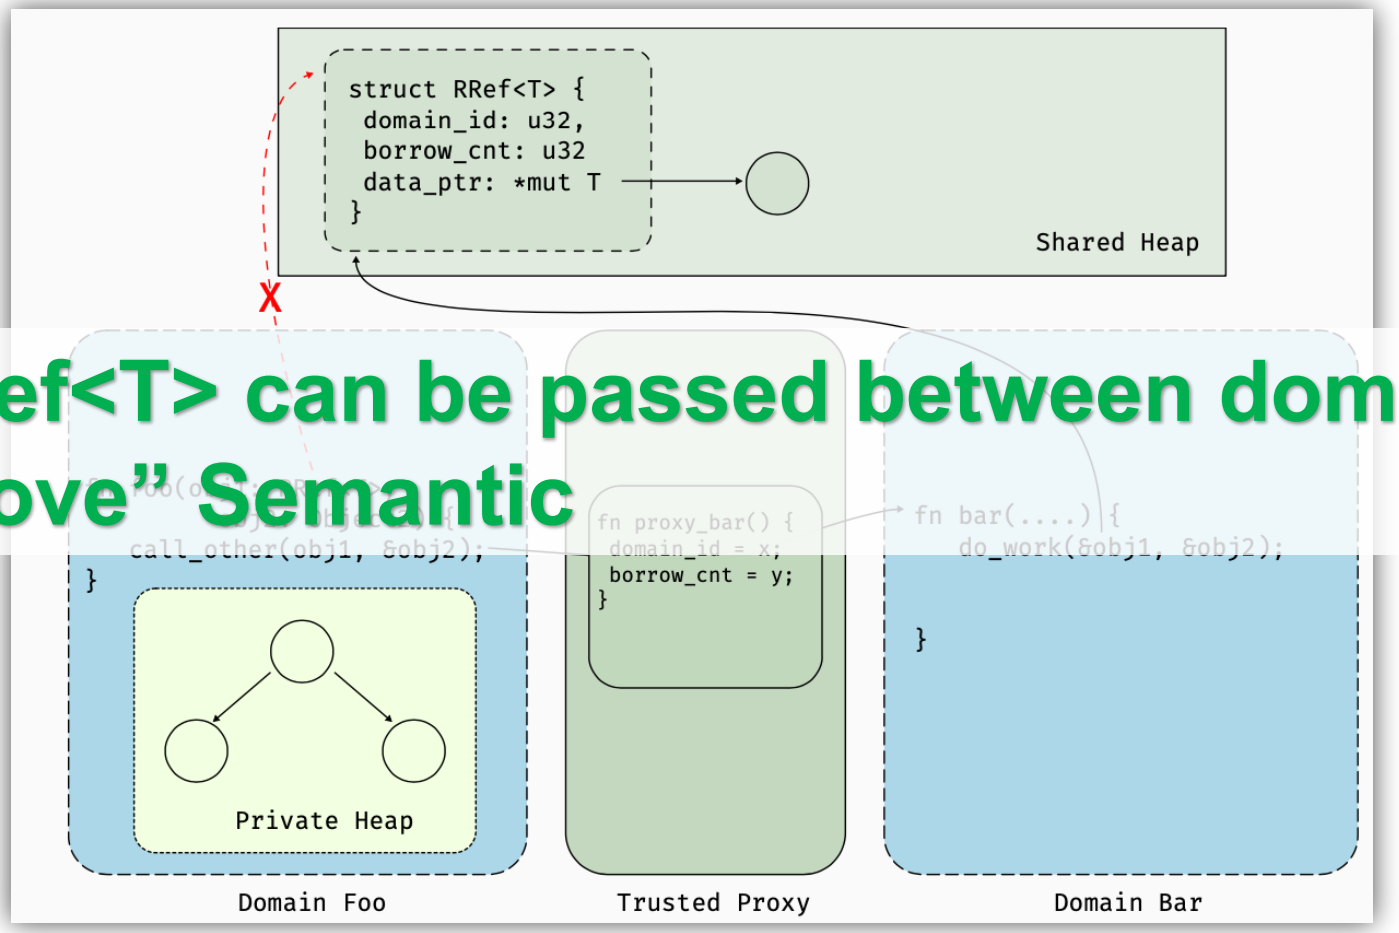

- RRef<T> can be passed between domains
- "Move" Semantic

3.6 Ownership Tracking

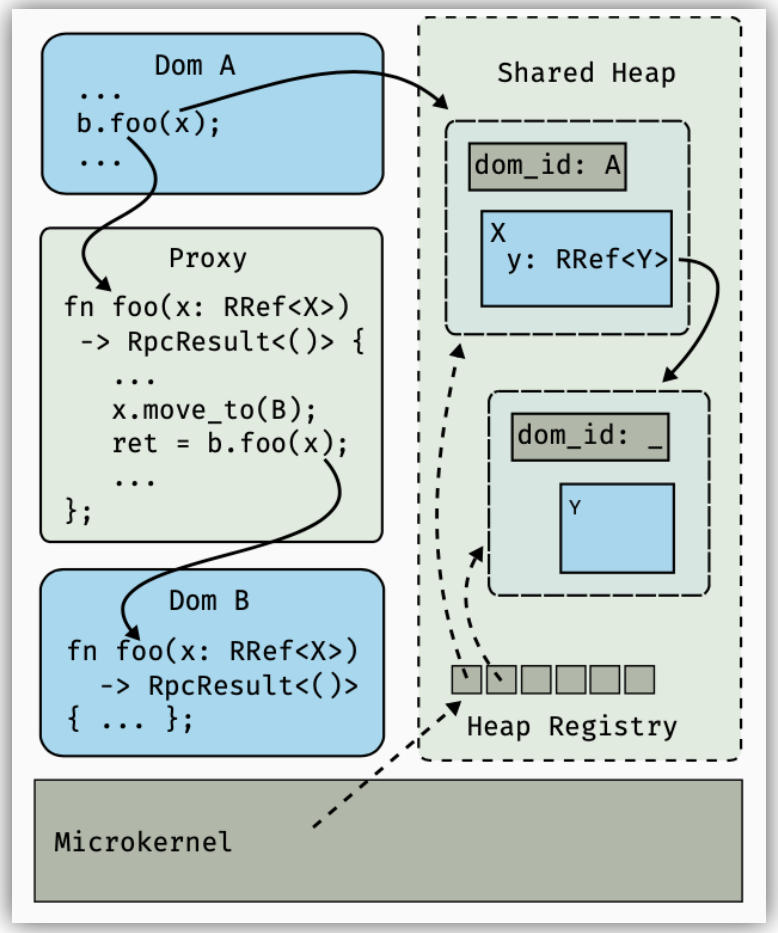

3.6 Ownership Tracking

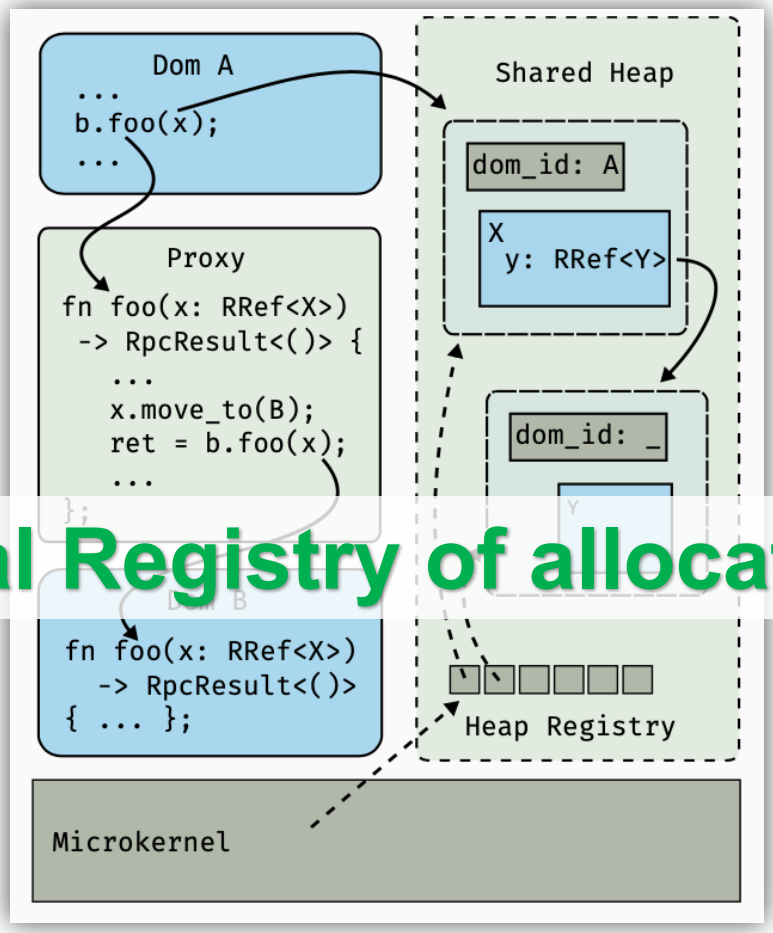

Global Registry of allocated objects

3.6 Ownership Tracking

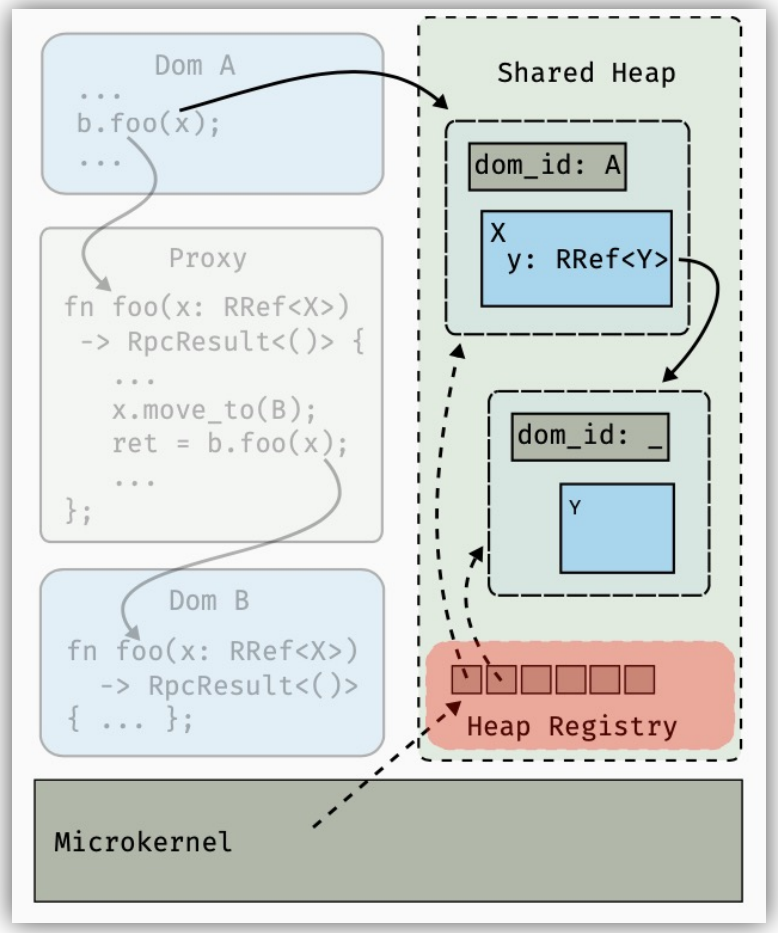

3.7 Cross-domain Call Proxying

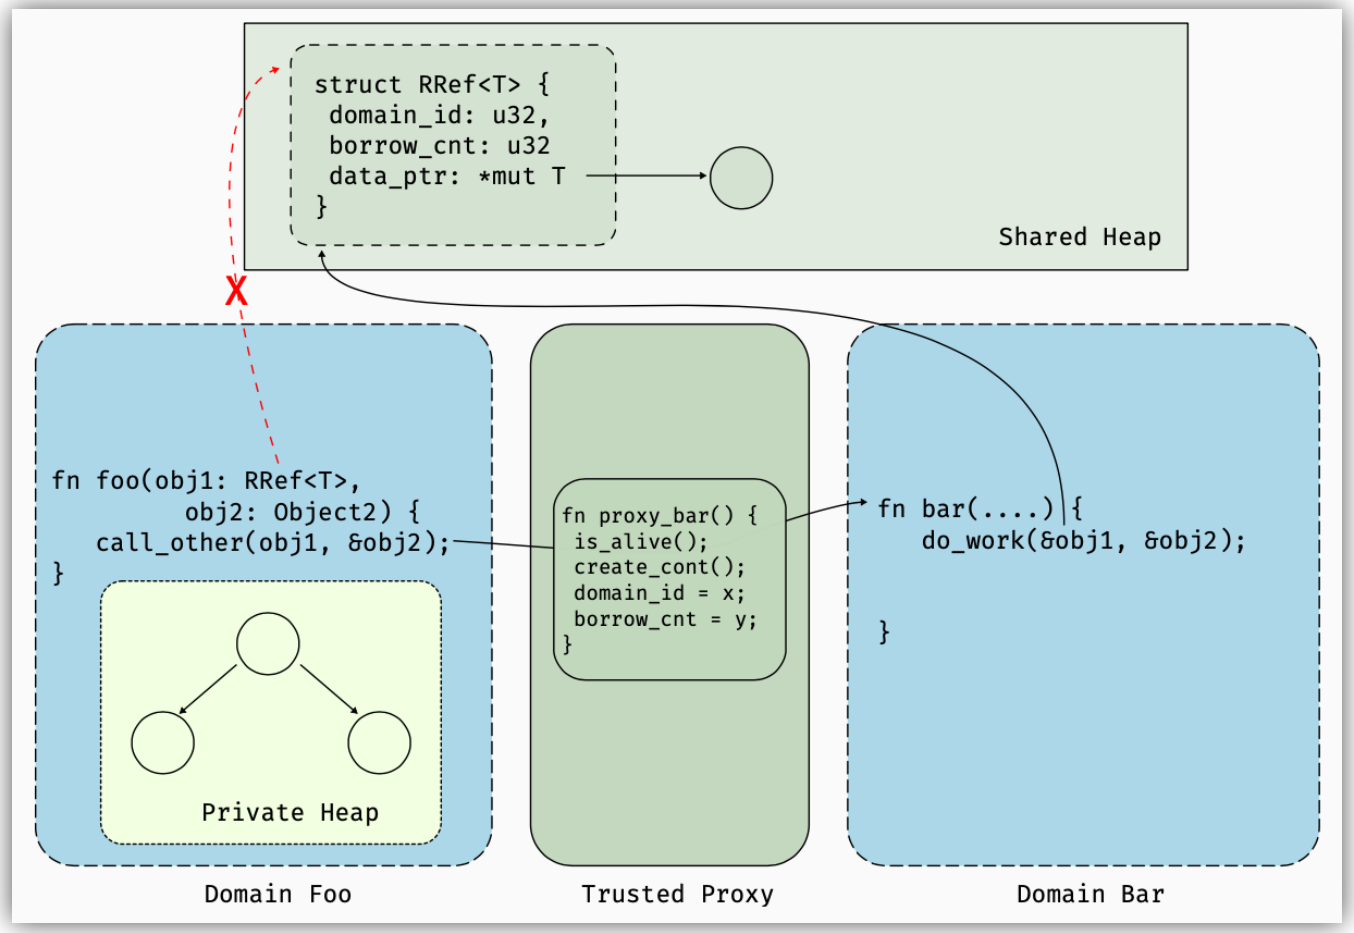

3.7 Cross-domain Call Proxying

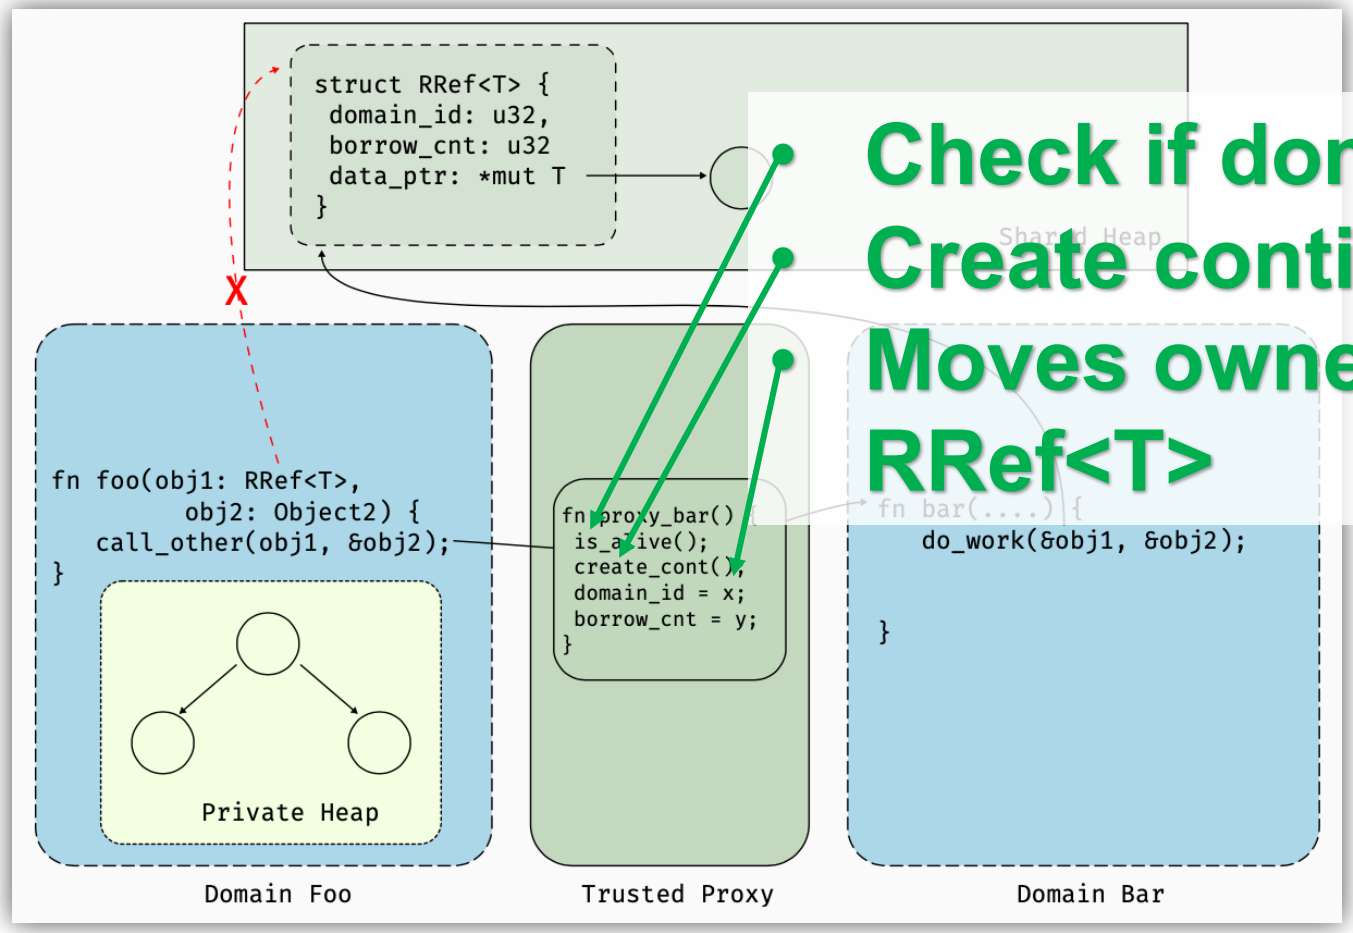

3.7 Cross-domain Call Proxying

- Where does the proxy comes from ?
- IDL (Interface Definition Language)

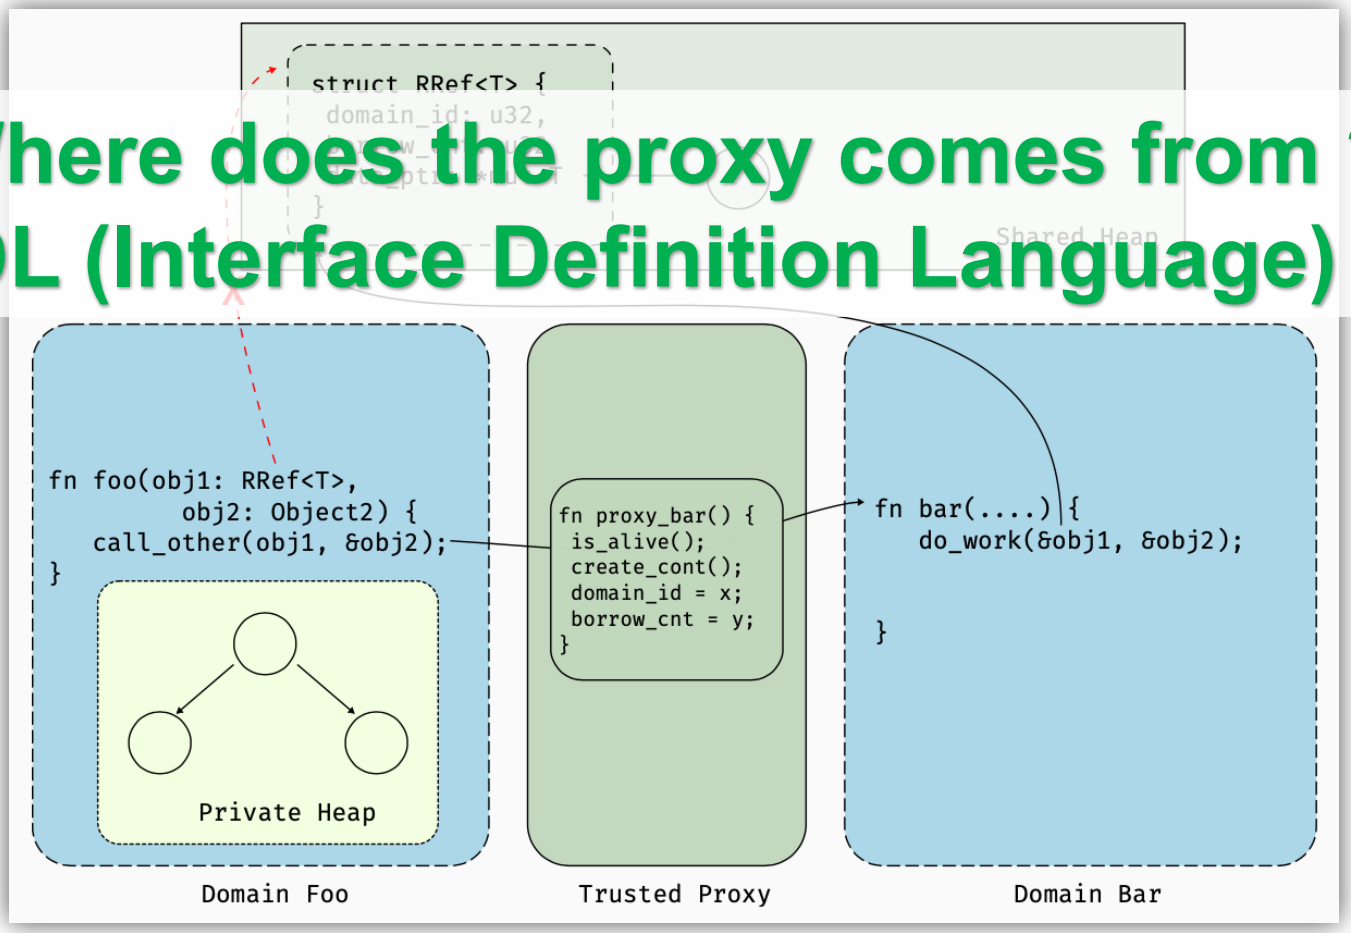

3.8 IDL & IDL Compiler

Example: Block Device Domain Interface

```
pub trait BDev {
    fn read(&self, block: u32, data: RRef<[u8; BSIZE]>)
        -> RpcResult<RRef<[u8; BSIZE]>>;
    fn write(&self, block: u32, data: &RRef<[u8; BSIZE]>)
        -> RpcResult<()>;
}

#[create]
pub trait CreateBDev {
    fn create(&self, pci: Box<dyn PCI>)
        -> RpcResult<(Box<dyn Domain>, Box<dyn BDev>)>
}
```

3.8 IDL & IDL Compiler

Interface

```
pub trait BDev {
    fn read(&self, block: u32, data: RRef<[u8; BSIZE]>)
        -> RpcResult<RRef<[u8; BSIZE]>>;
    fn write(&self, block: u32, data: &RRef<[u8; BSIZE]>)
        -> RpcResult<()>;
}

#[create]
pub trait CreateBDev {
    fn create(&self, pci: Box<dyn PCI>)
        -> RpcResult<(Box<dyn Domain>, Box<dyn BDev>)>
}
```

Create a domain

3.8 IDL & IDL Compiler

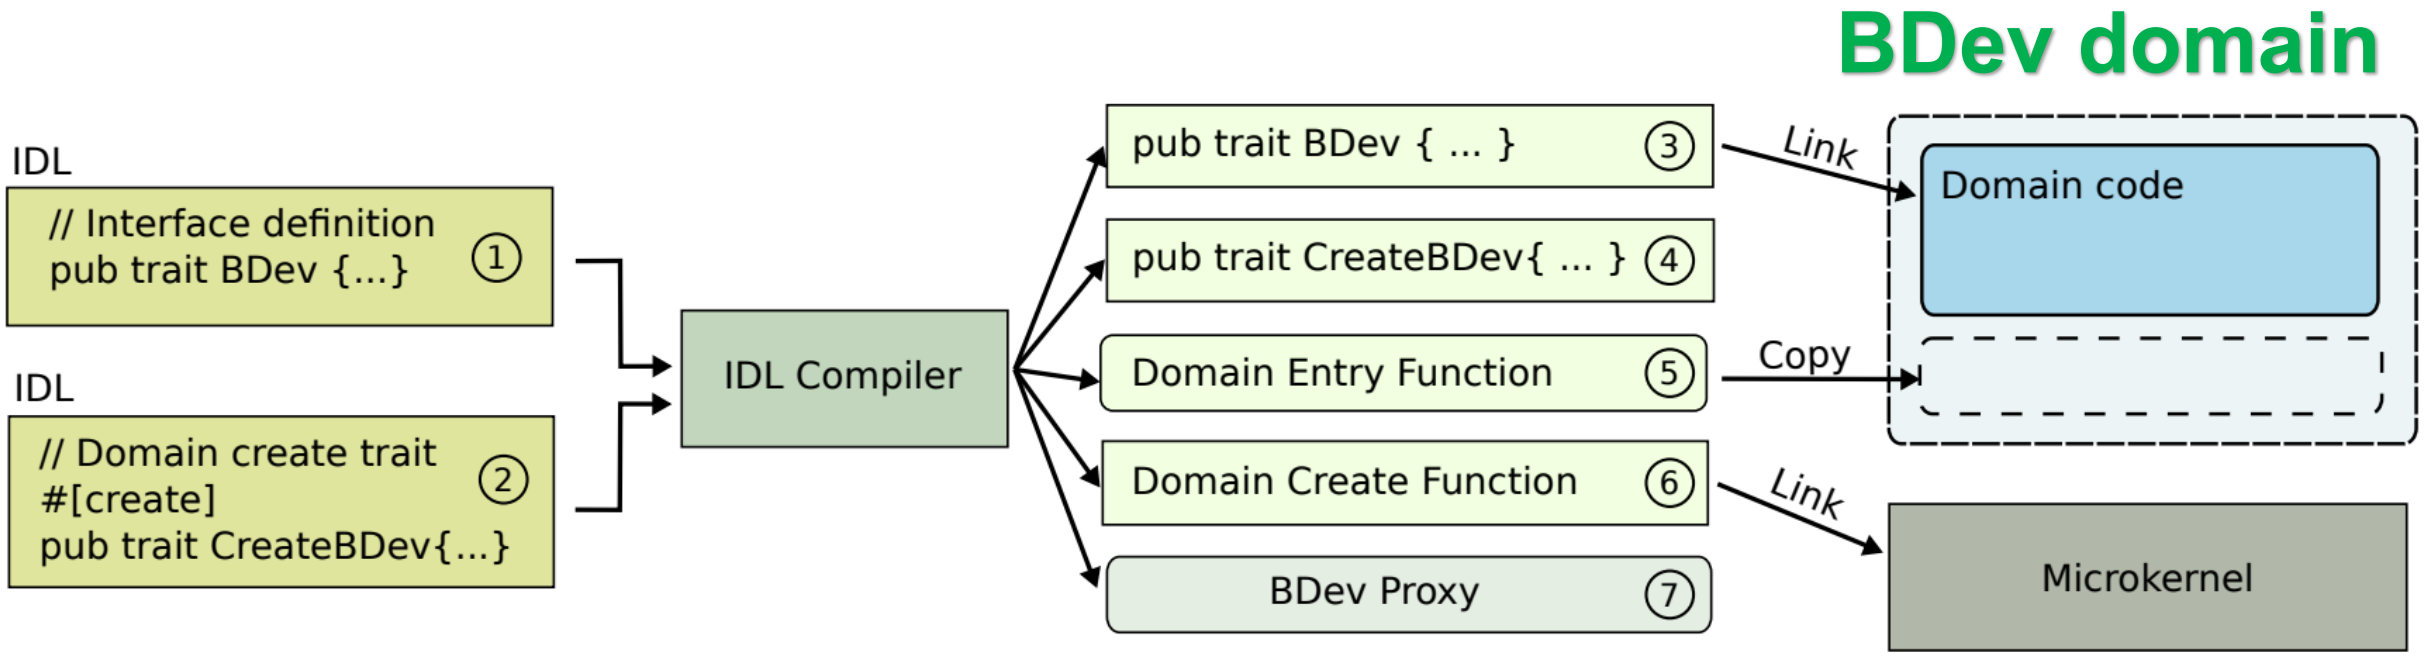

3.8 IDL & IDL Compiler

Static Analysis on AST
to extract interface definition

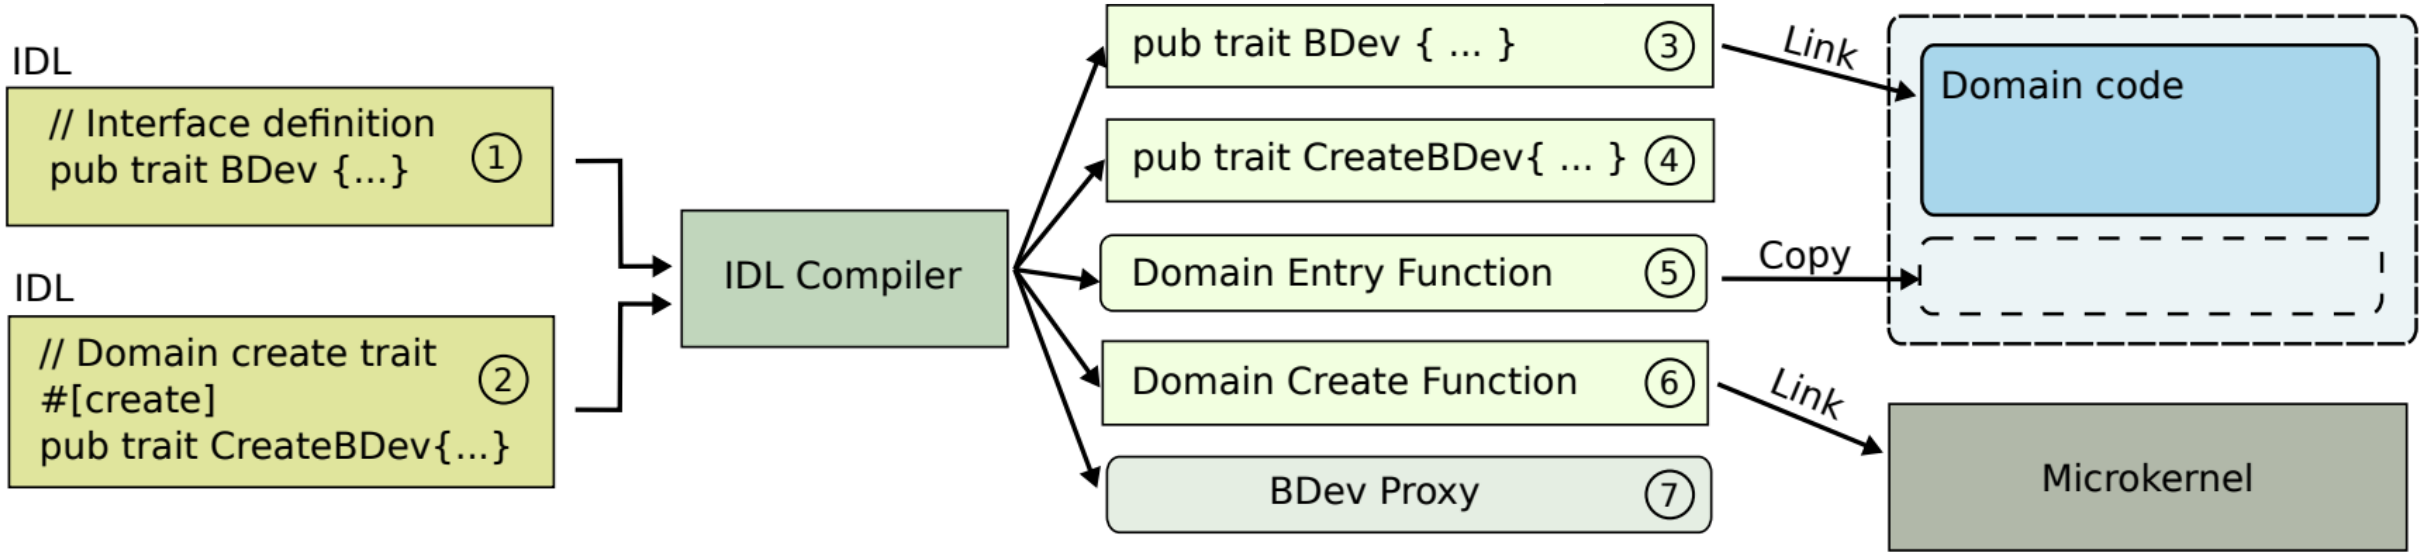

3.9 Summary of RedLeaf

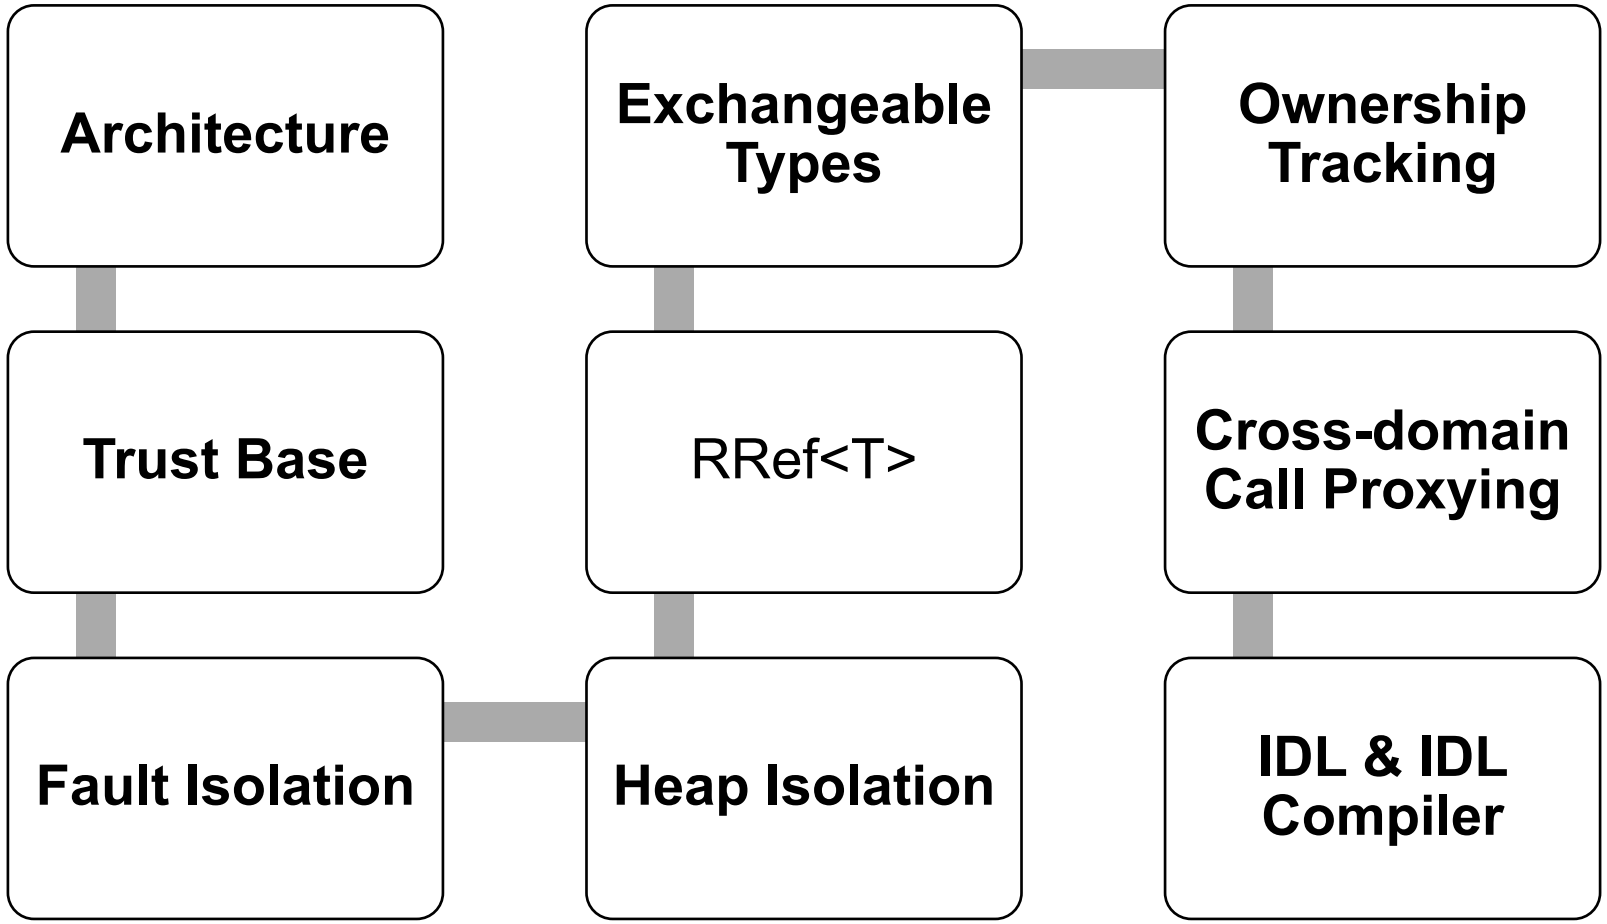

3.10 Device Driver Recovery

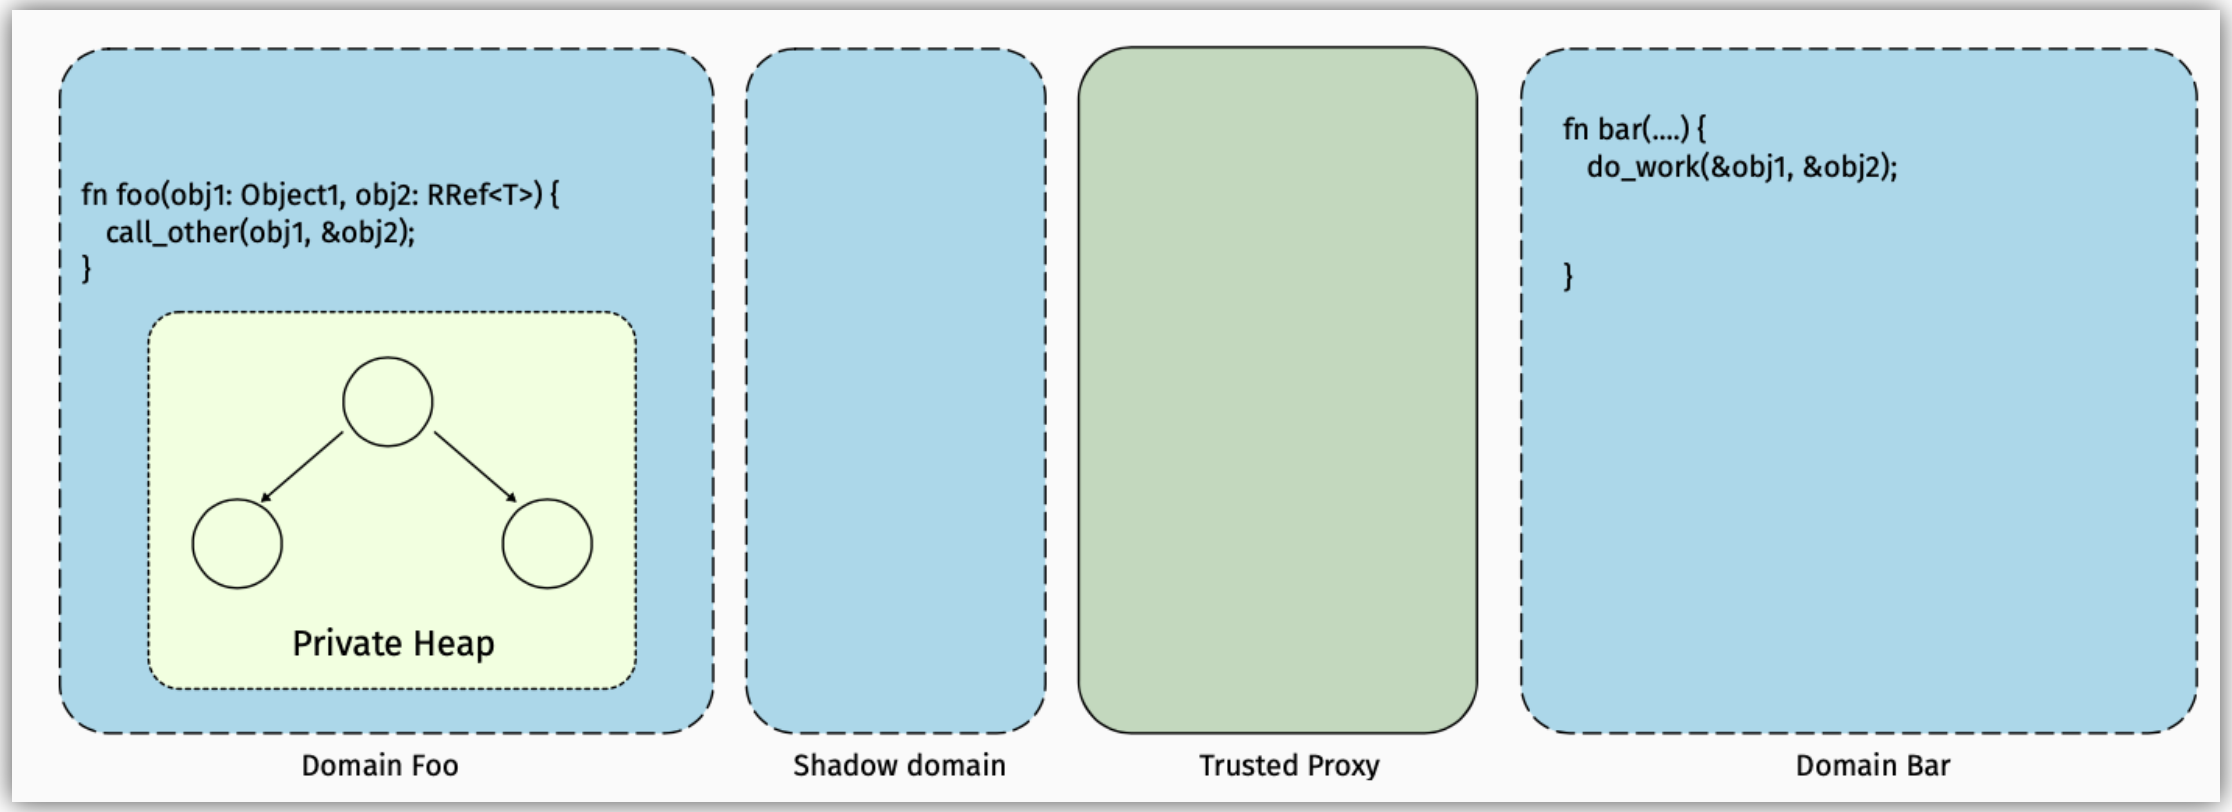

3.10 Device Driver Recovery

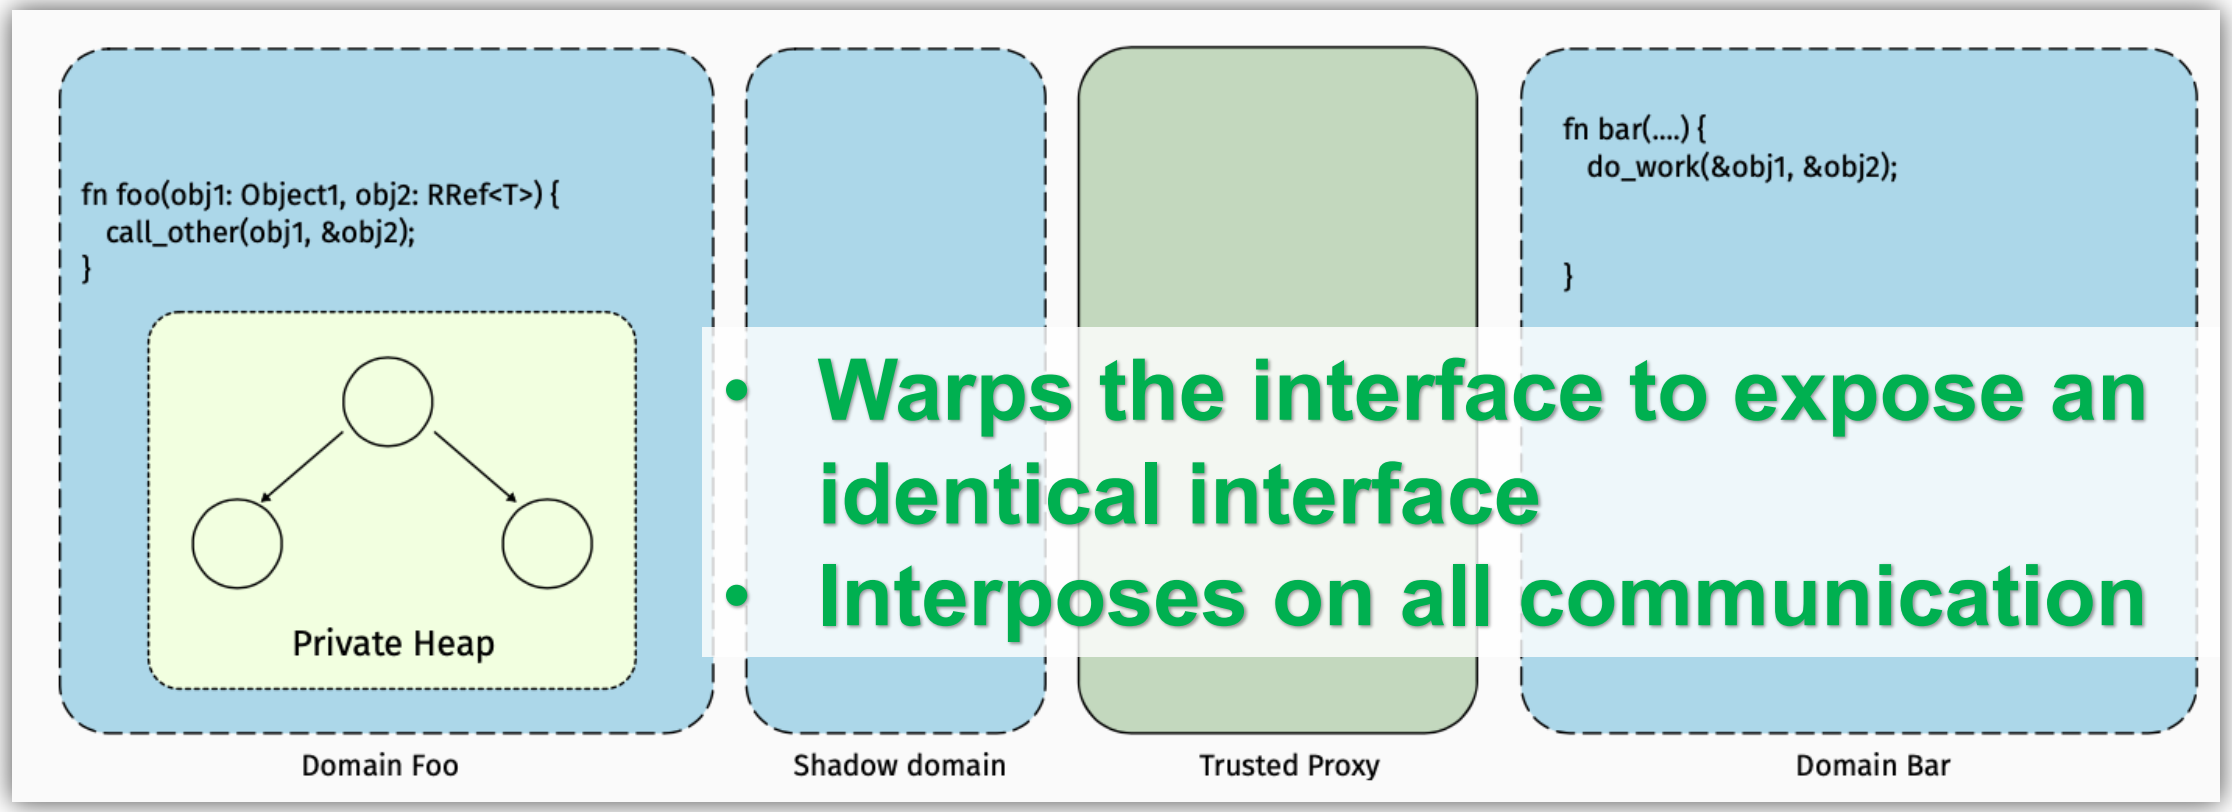

- Warps the interface to expose an identical interface
- Interposes on all communication

3.10 Device Driver Recovery

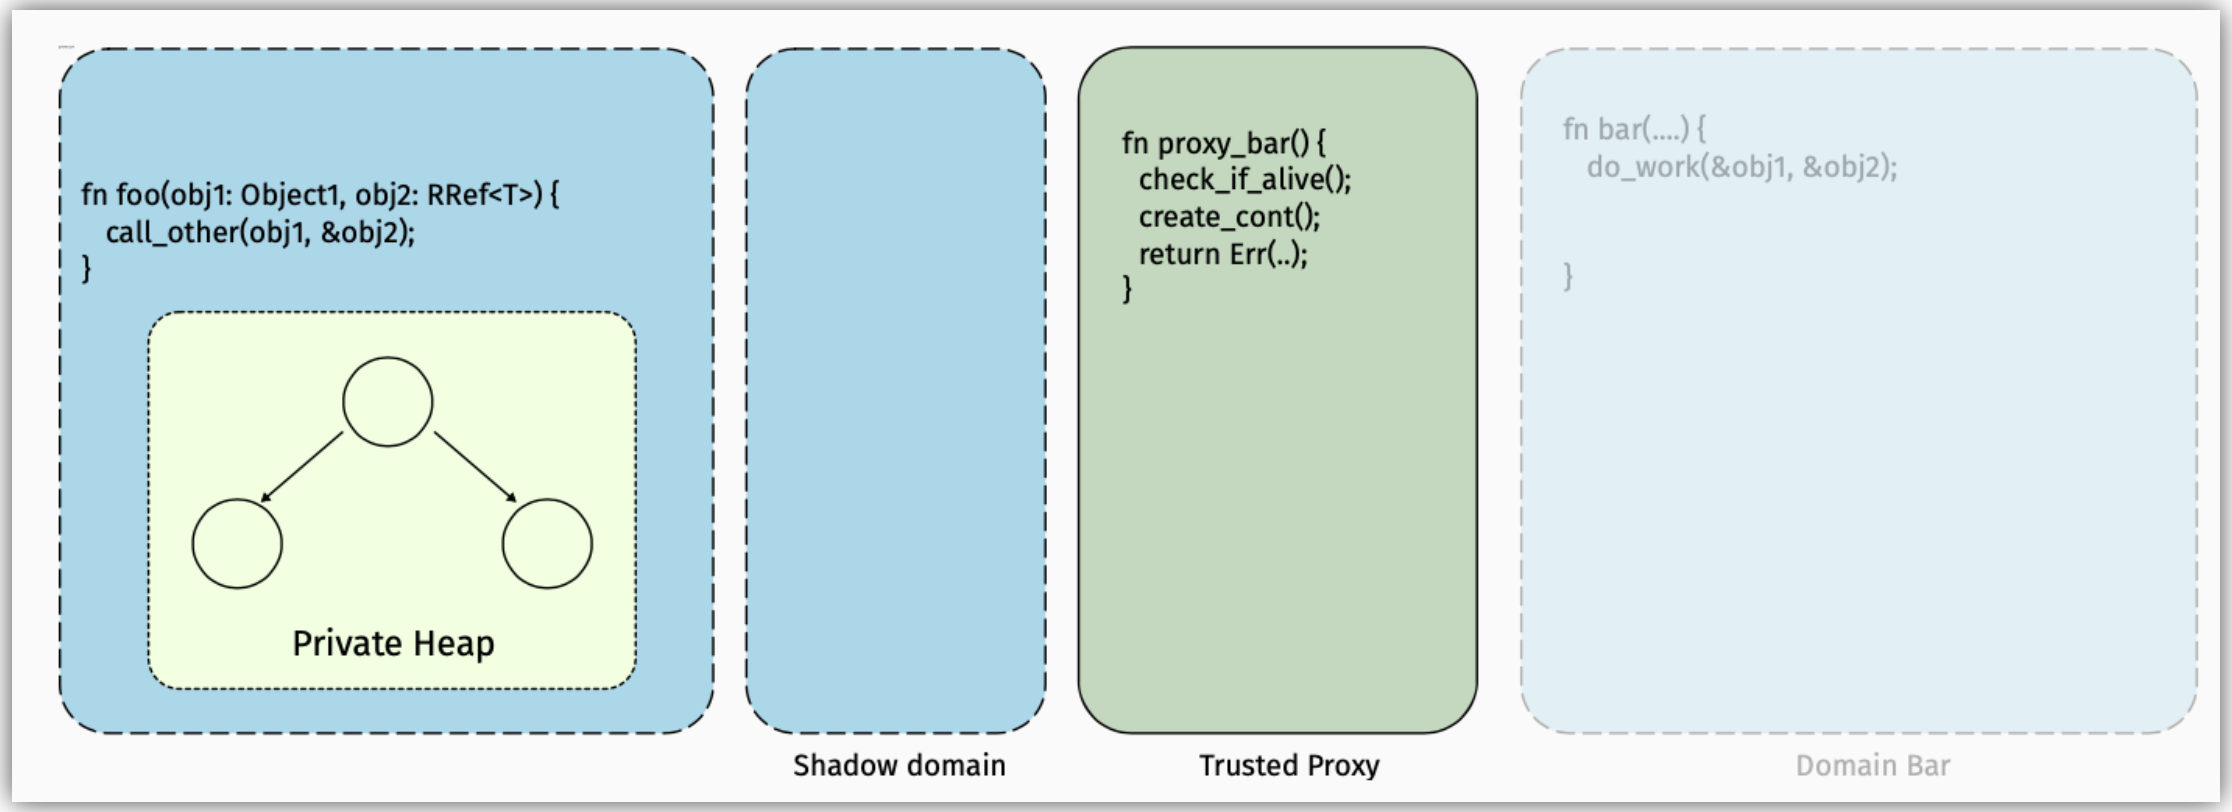

3.10 Device Driver Recovery

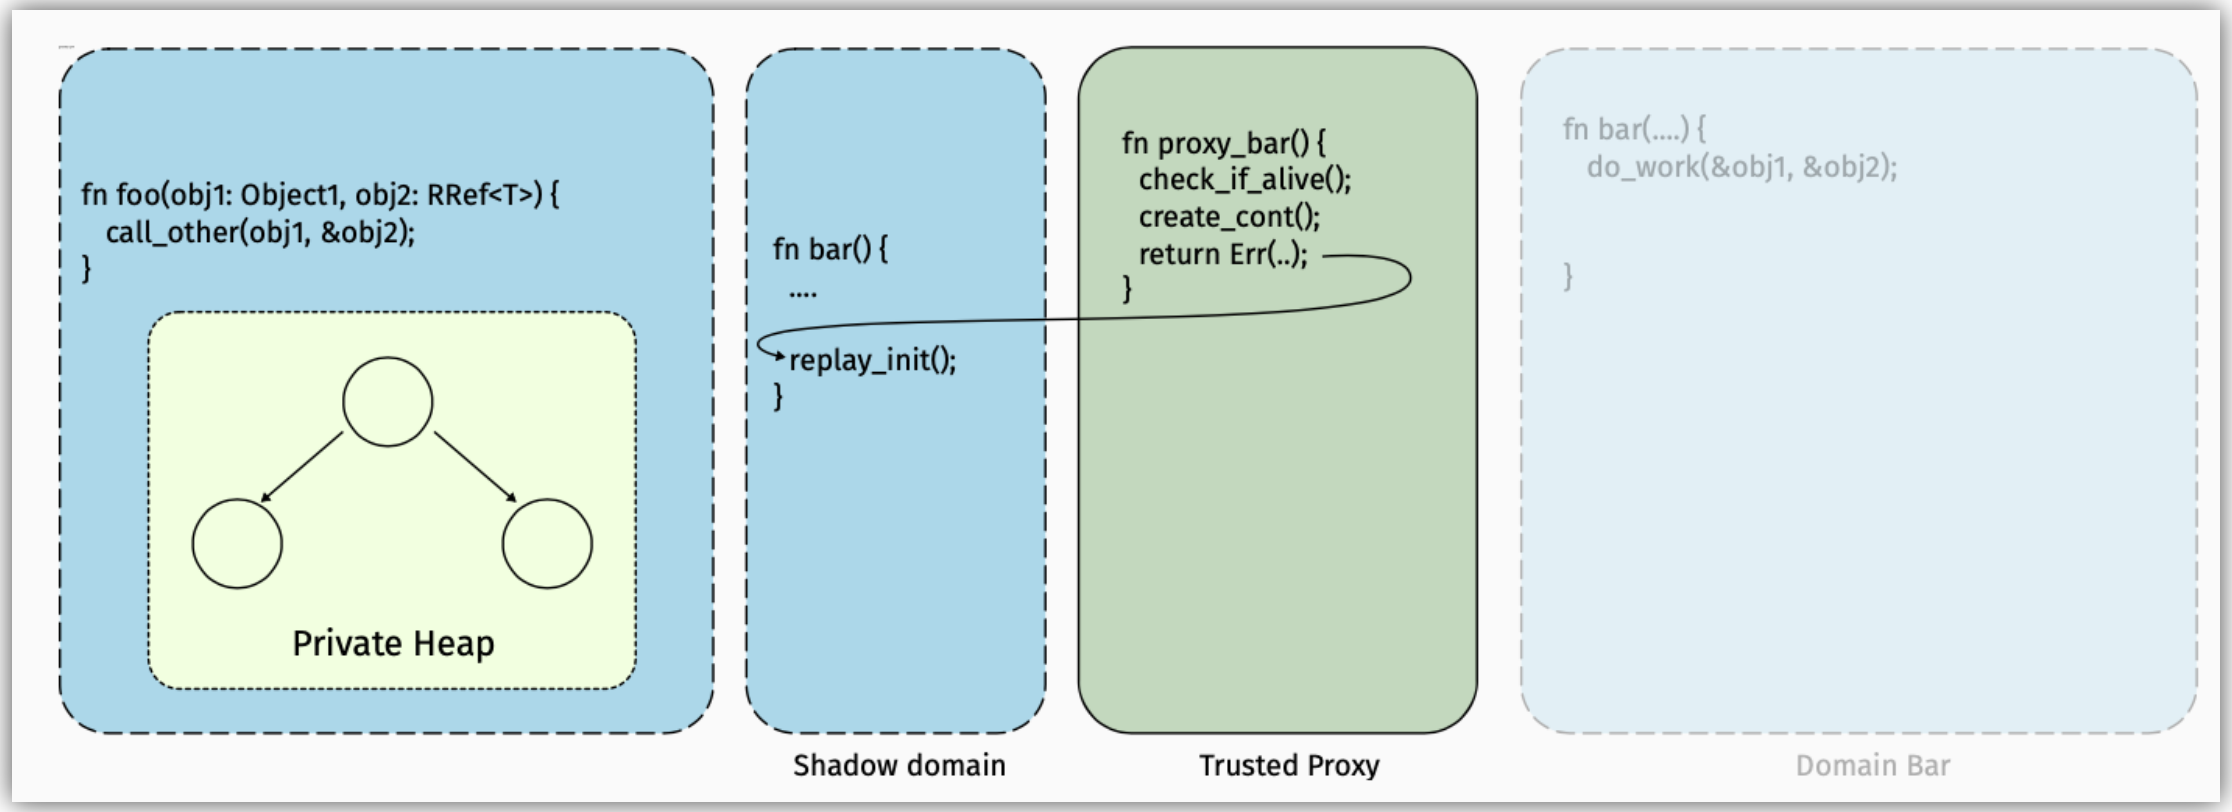

3.10 Device Driver Recovery

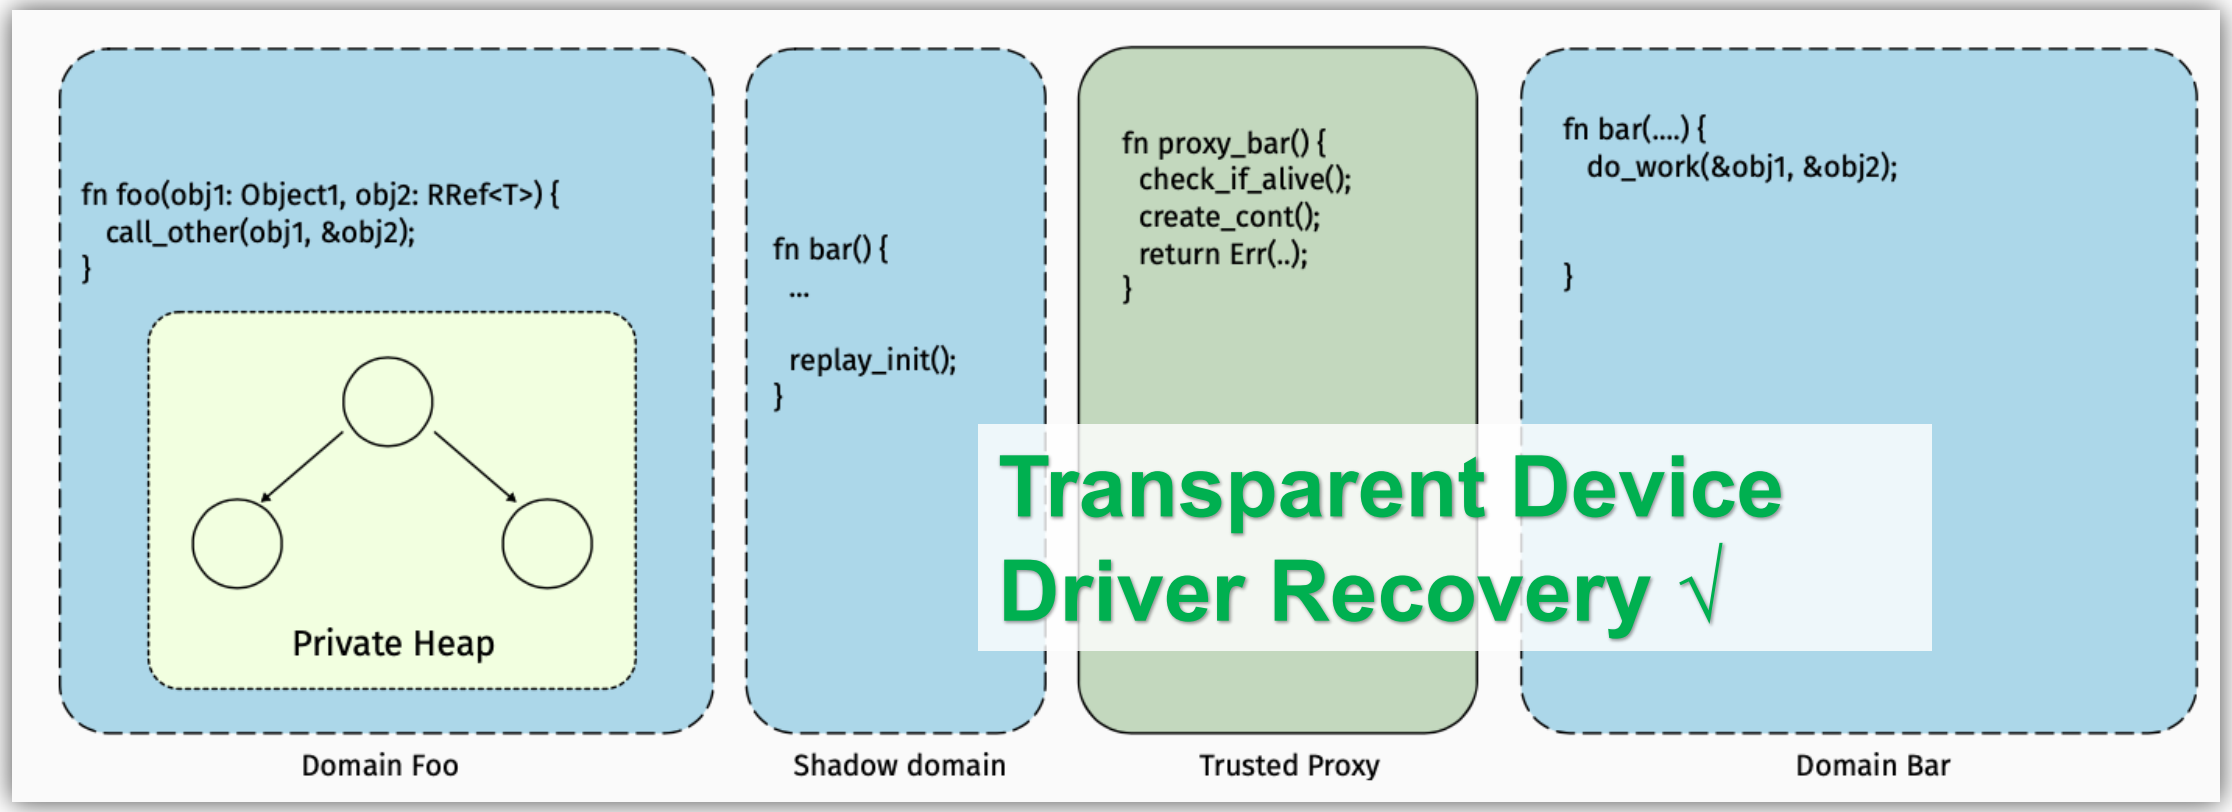

4. Evaluation

4.1 Communication Cost

Operation	Cycles
seL4	834
VMFUNC	169
VMFUNC-based call/reply invocation	396
RedLeaf cross-domain invocation	124
RedLeaf cross-domain invocation (passing an `RRef<T>`)	141
RedLeaf cross-domain invocation via shadow	279
RedLeaf cross-domain via shadow (passing an `RRef<T>`)	297

4.2 Language Overhead

- Hashtable - (FNV hash, open addressing, <8B, 8B>)
- C, Idiomatic Rust, C-style Rust,
- C-style Rust: No higher order functions **usize, usize**
- Idiomatic Rust - **Option<(usize, usize)>**
- Vary the size (2^{12} to 2^{26} at 75% full)

4.2 Language Overhead

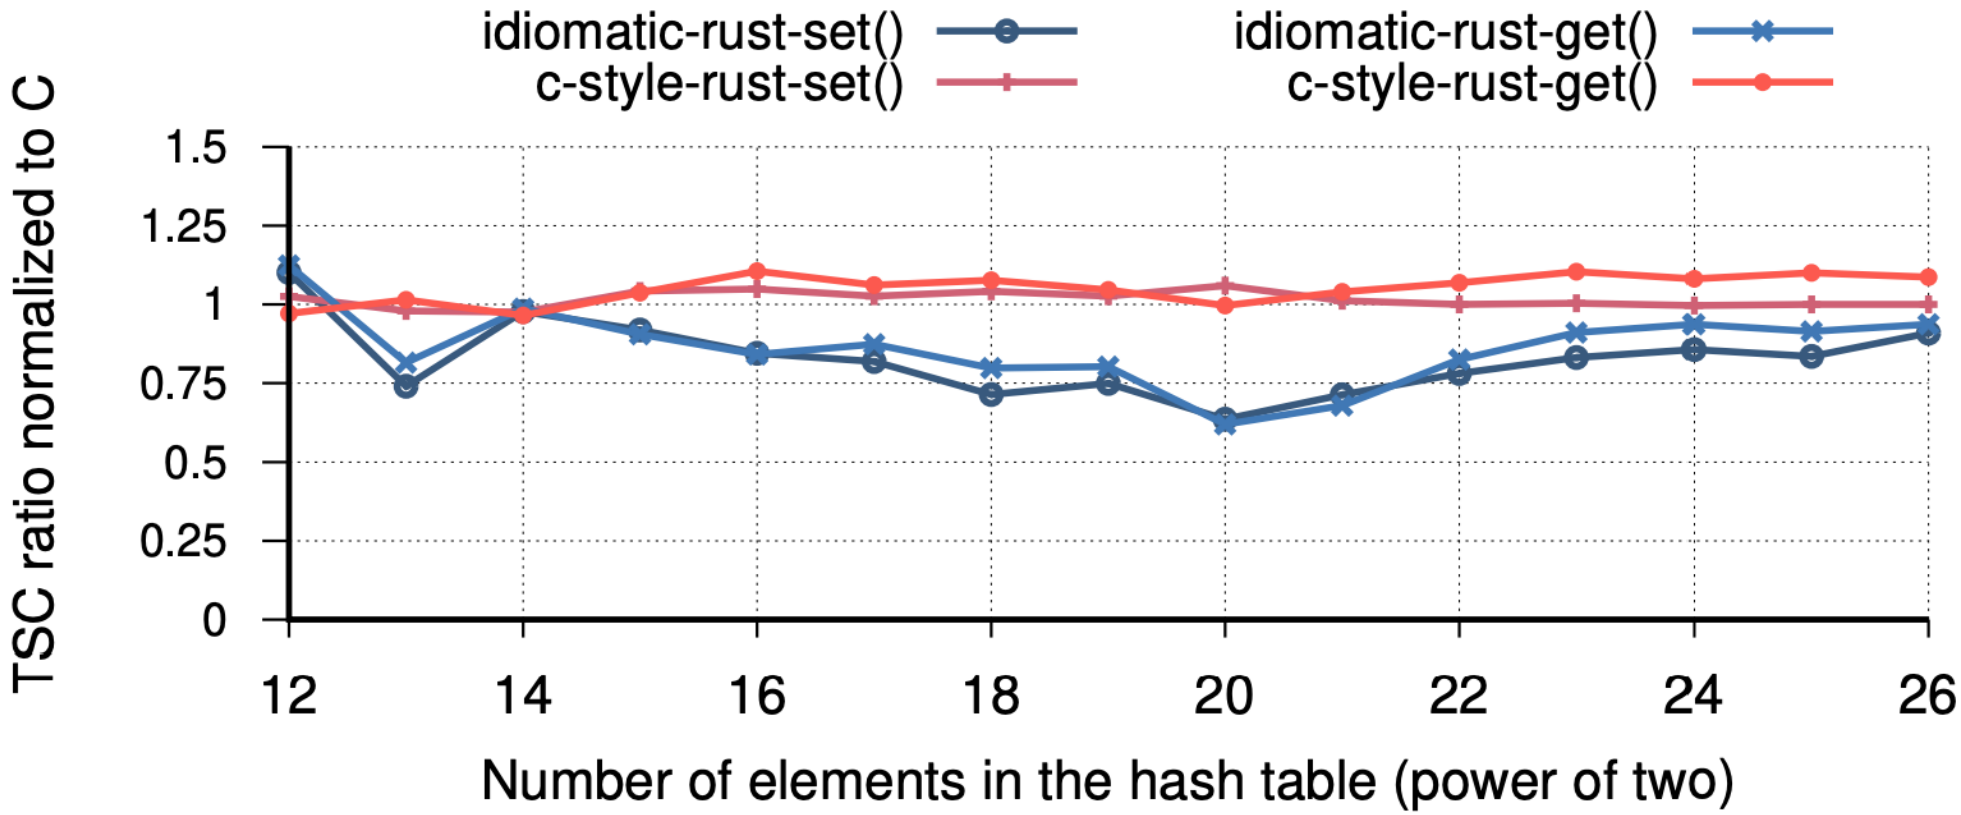

4.3 Device Drivers: ixgbe

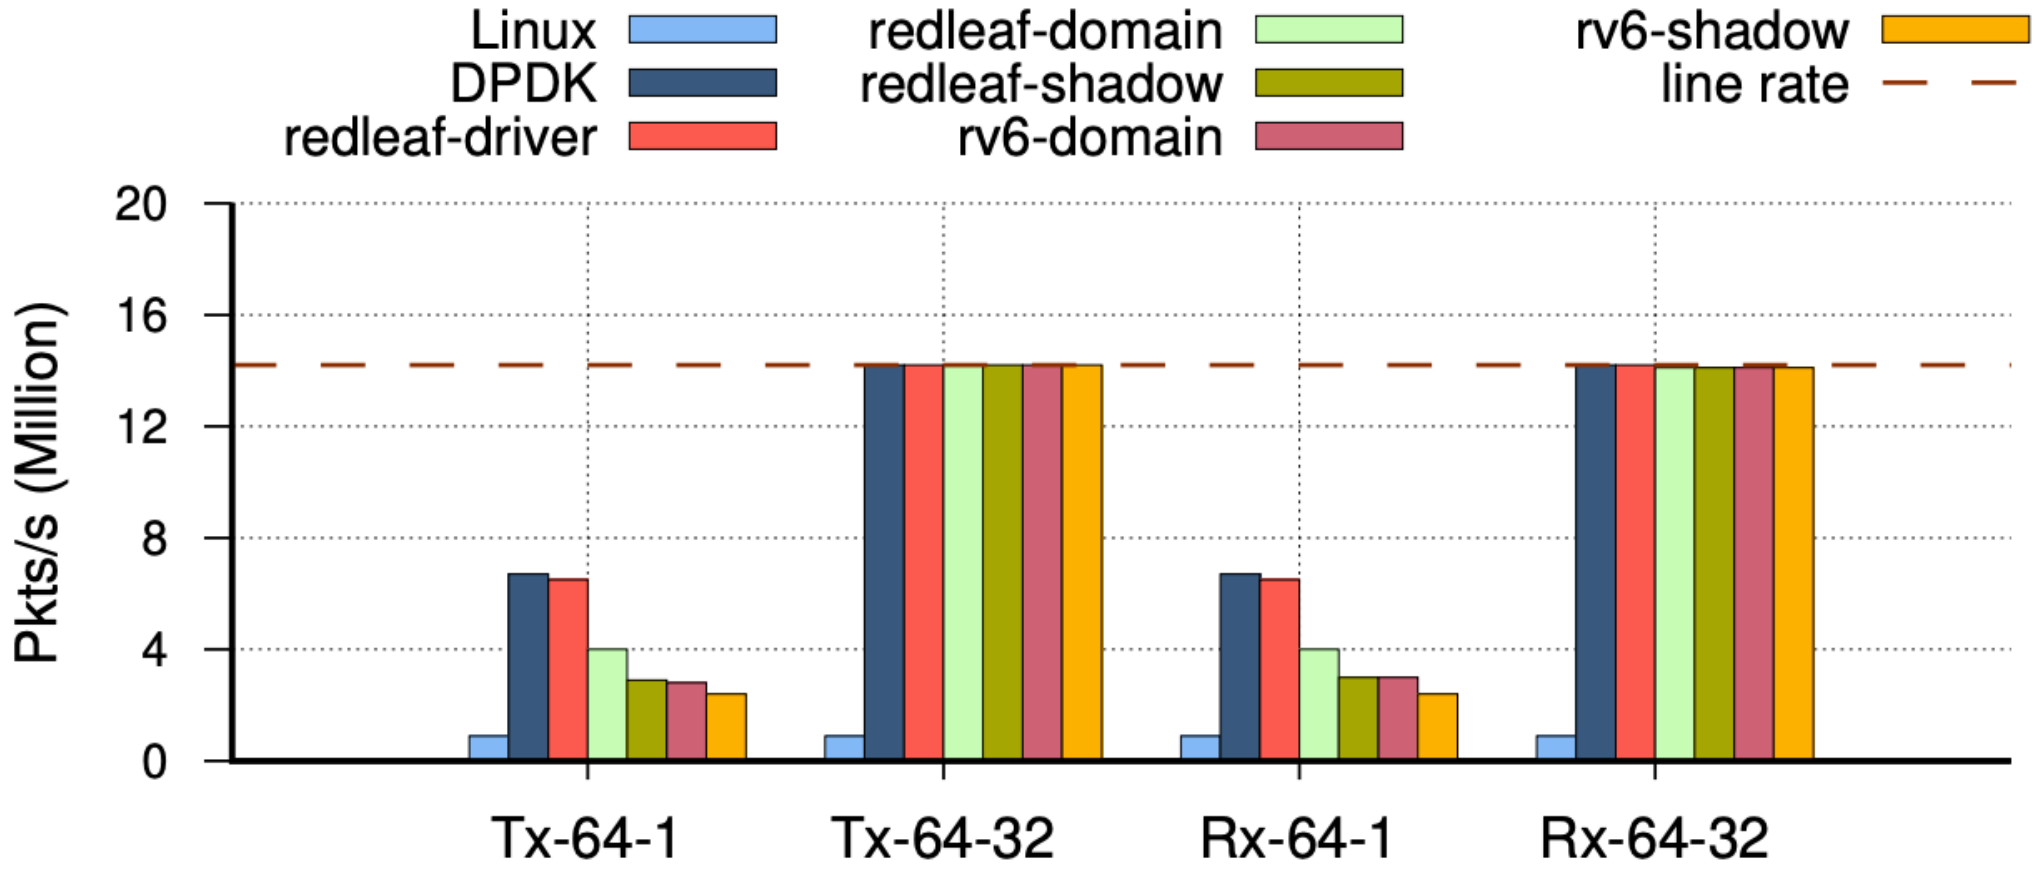

4.3 Device Drivers: NVMe

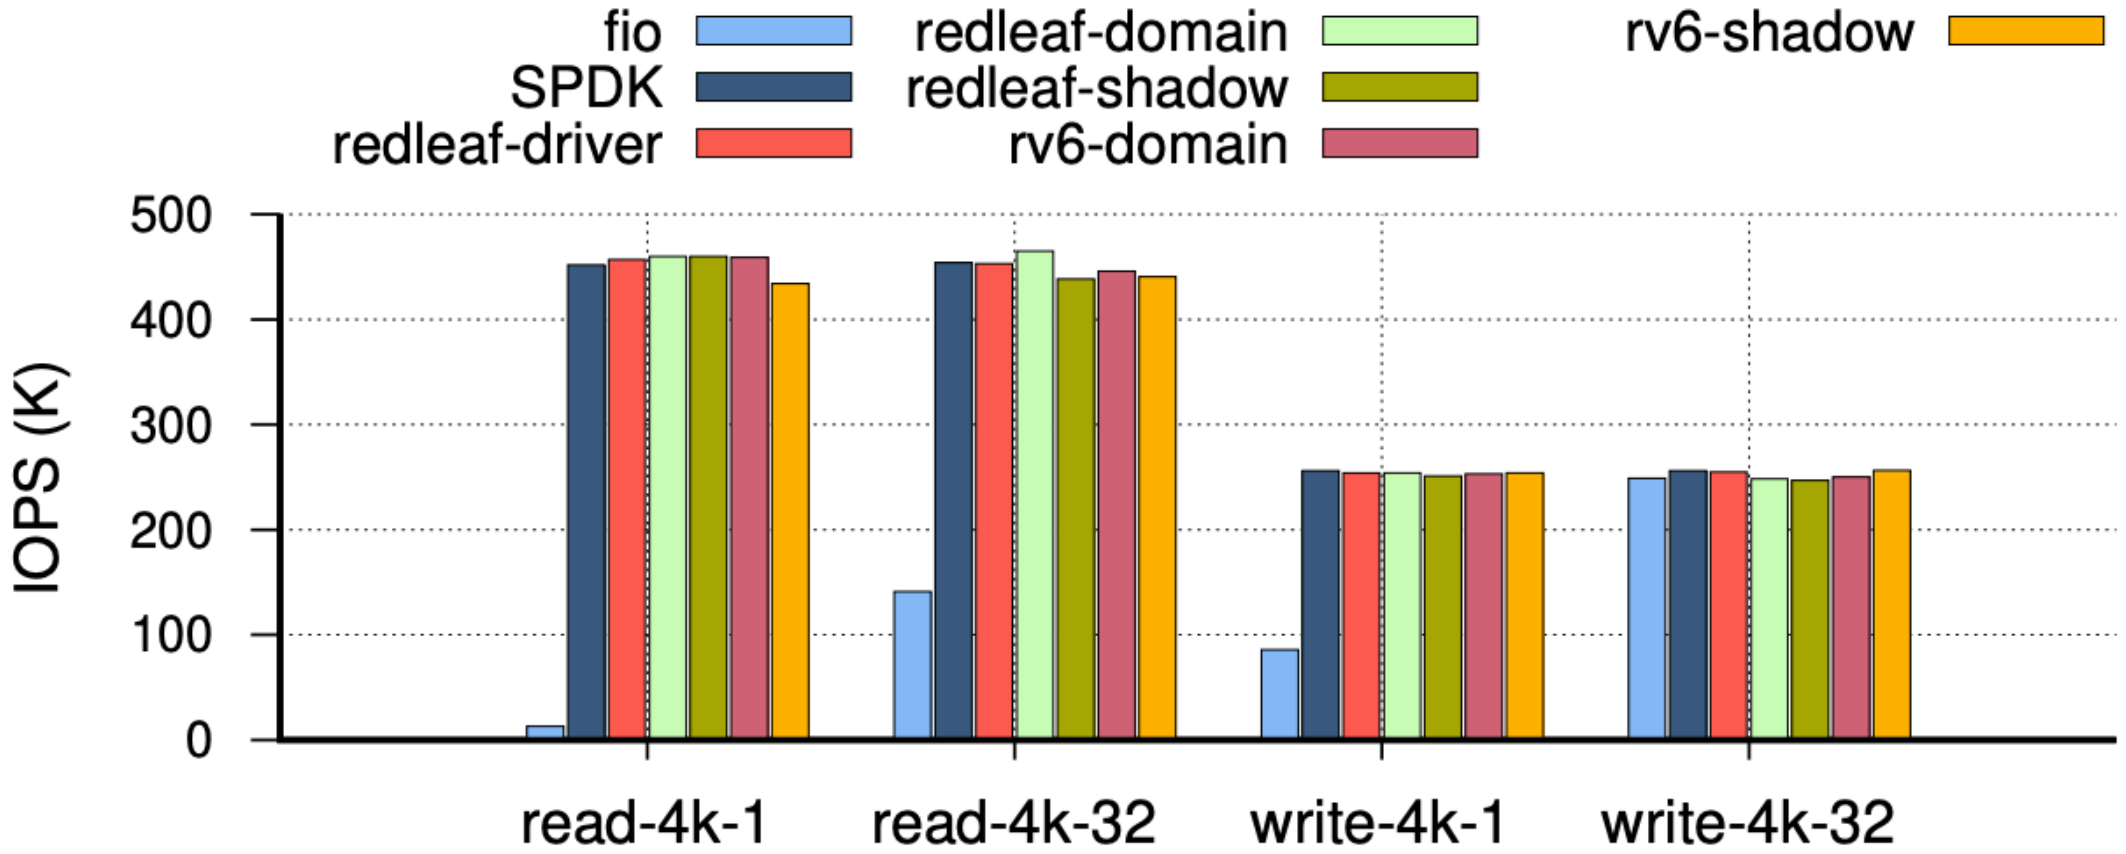

4.4 Application Test: Maglev

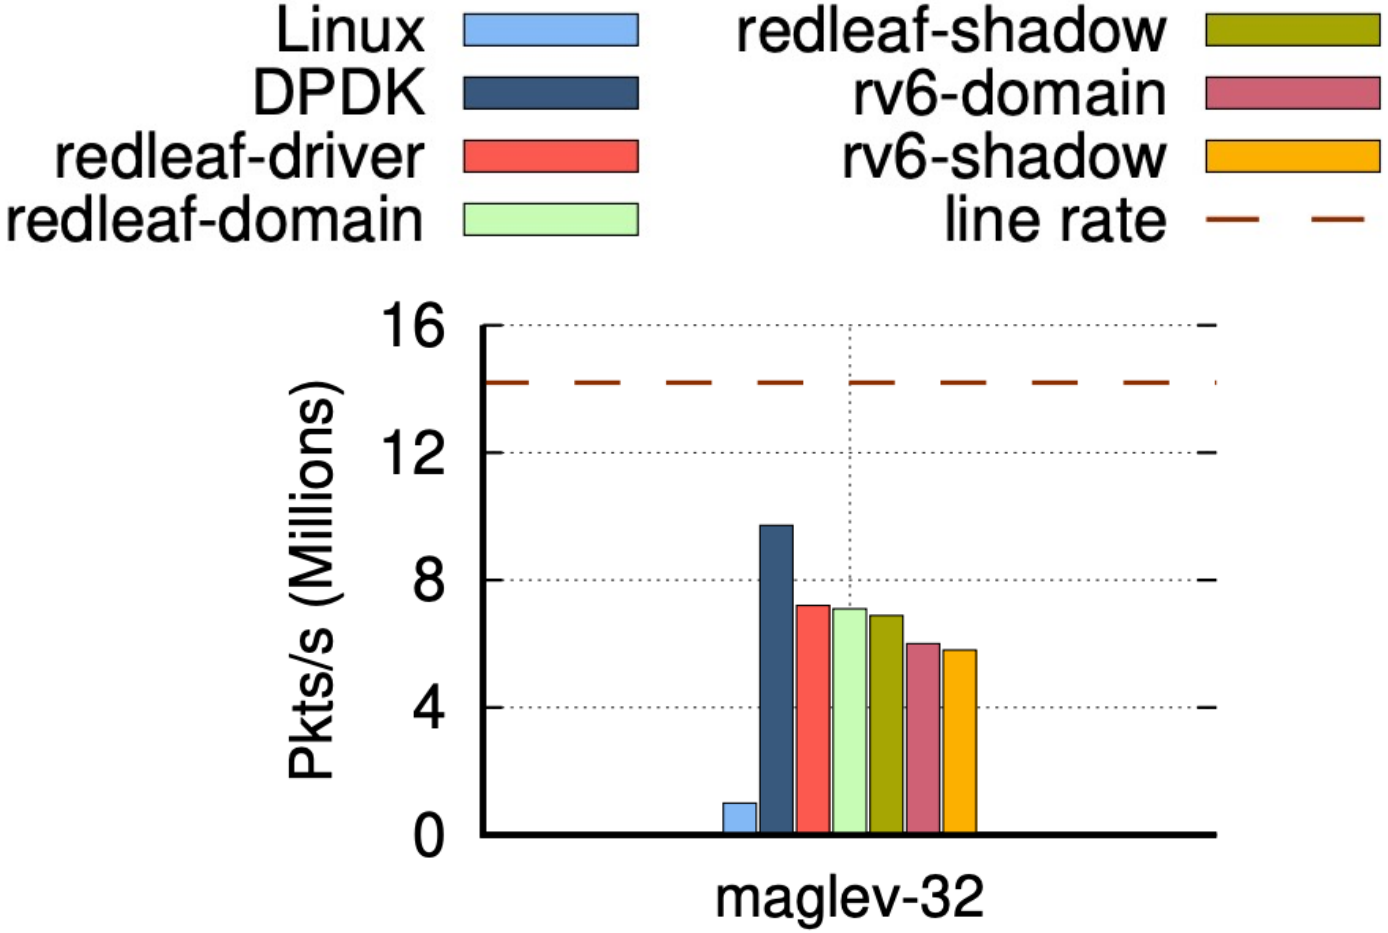

4.4 Application Test: Httpd

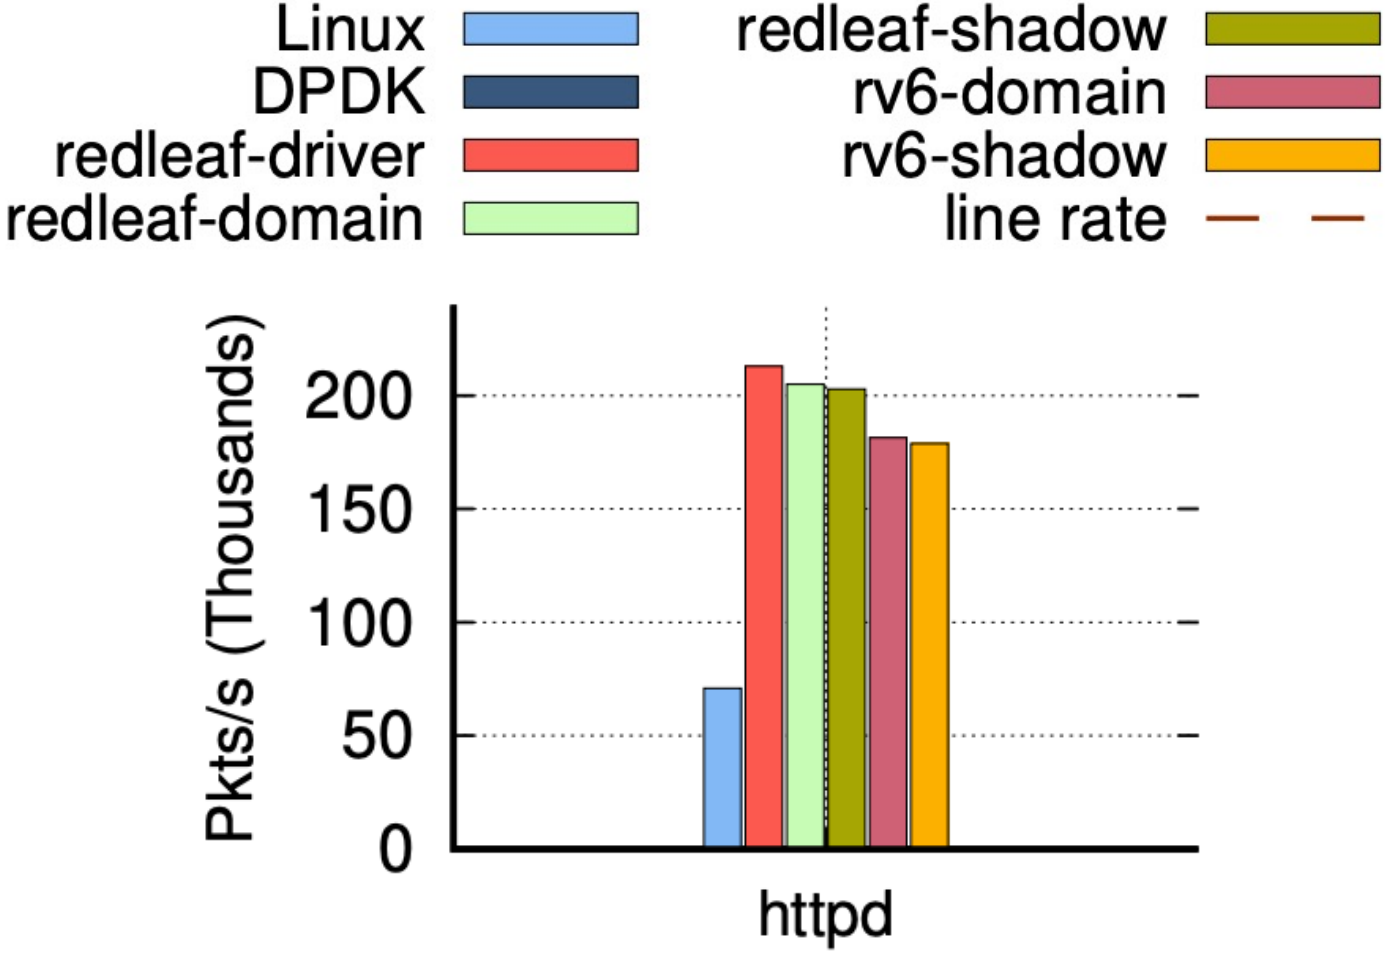

4.4 Application Test: KV-Store

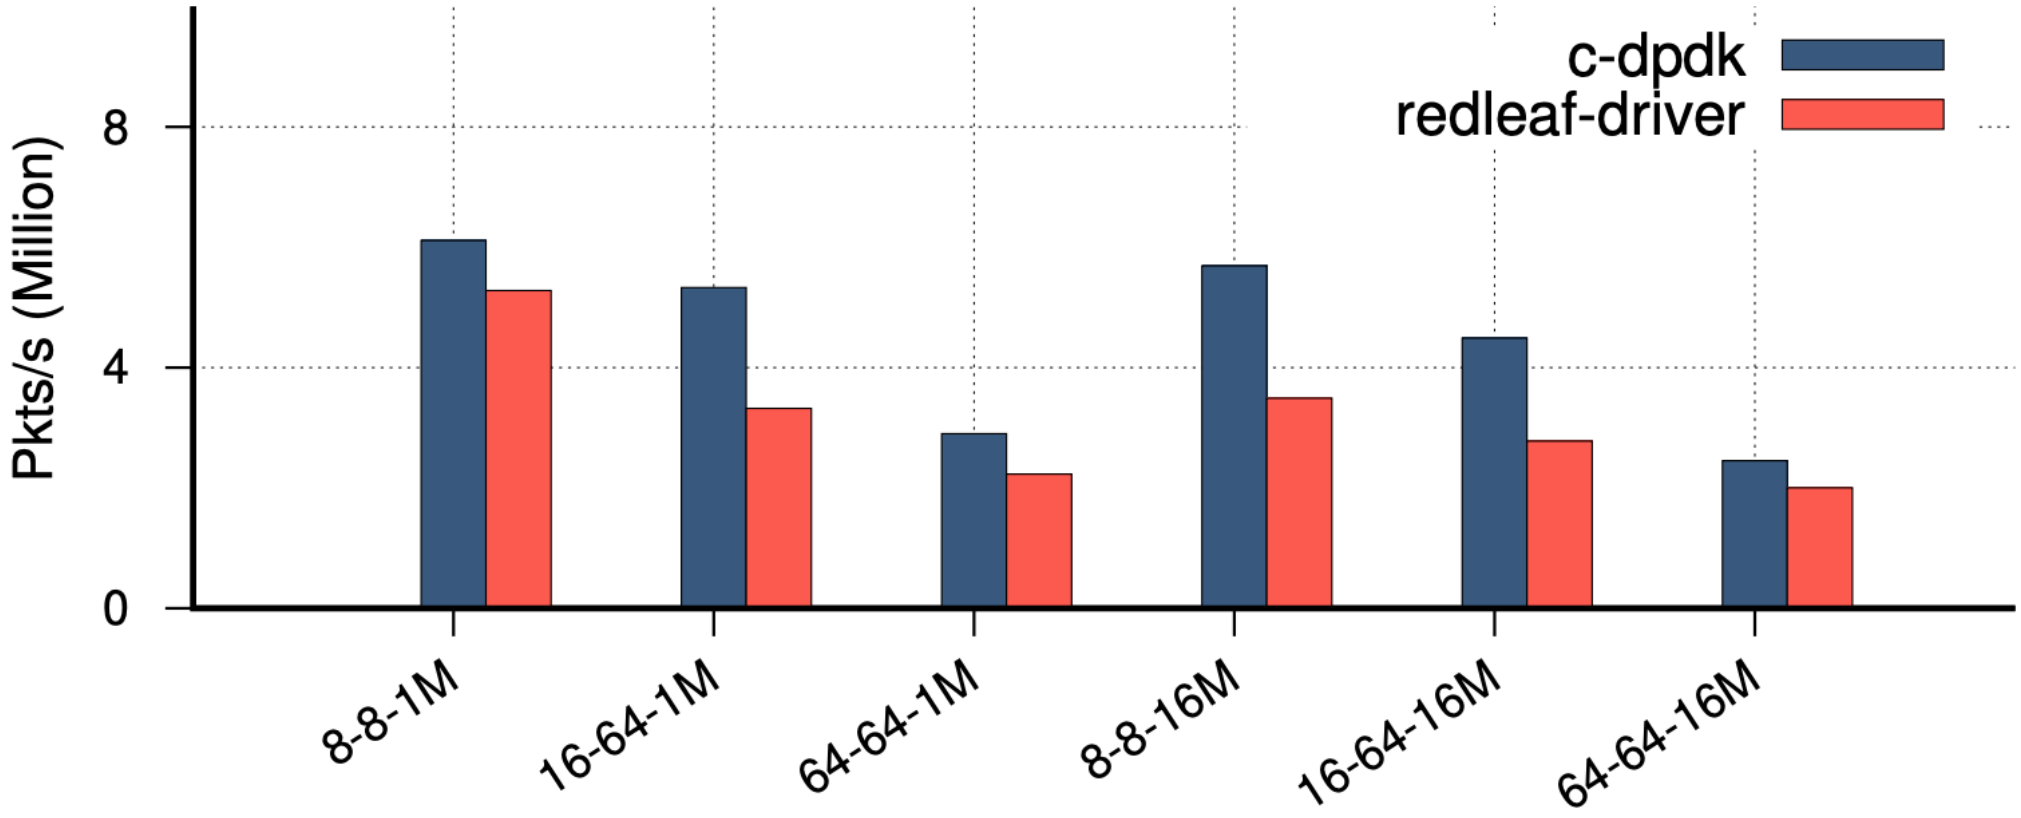

4.5 Device Driver Recovery

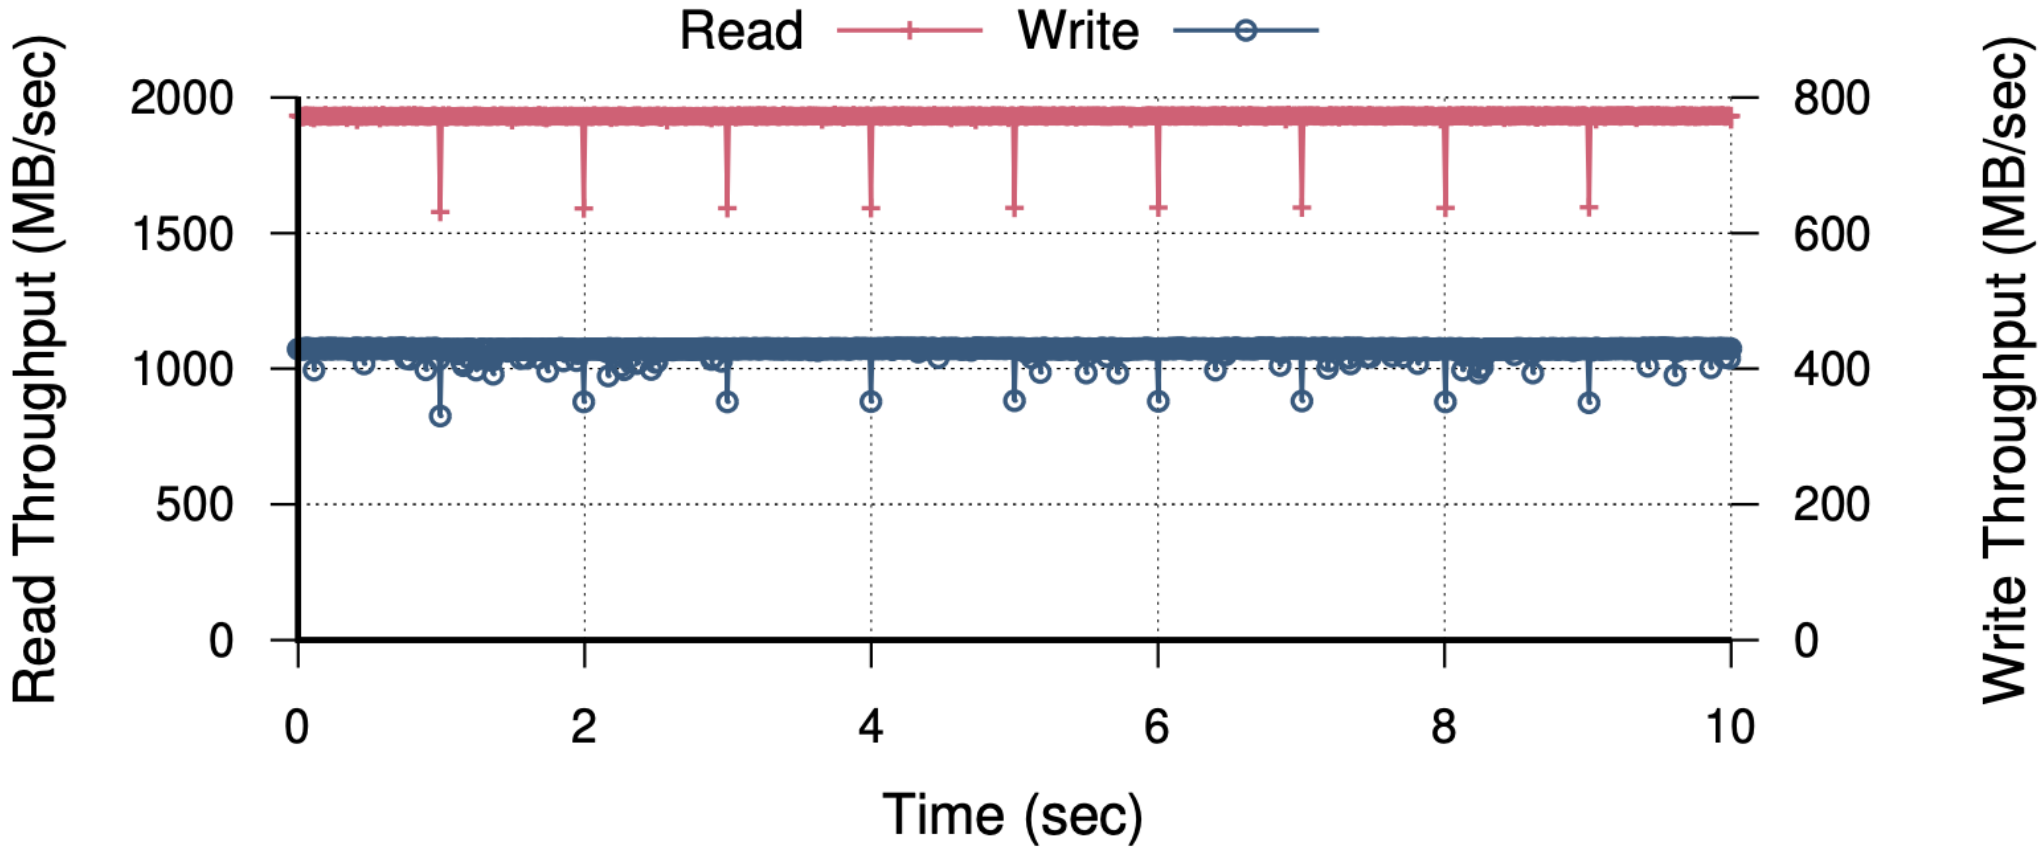

5. Conclusion & Insight

5.1 Conclusion

- Heap isolation, exchangeable types, ownership tracking, interface validation, cross-domain call proxying
- Provides a collection of mechanisms for enabling isolation
- A step forward in enabling future system architectures
  - Secure kernel extensions
  - fine-grained access control
  - transparent recovery