

Poster: A Systematic Review of Migration-Relevant Software Idioms and Their Interactions with Capability Systems

Fatih Durmaz

CISPA Helmholtz Center for Information Security
fatih.durmaz@cispa.de

Sven Bugiel

CISPA Helmholtz Center for Information Security
bugiel@cispa.de

Abstract—Software compartmentalization reduces vulnerabilities’ impact by decomposing software into isolated, least-privilege components. Capability systems support this goal by replacing ambient authority with explicit, delegable rights and fine-grained protection. However, legacy software was usually not written for such models, so migration requires simultaneous reasoning about correctness, security, performance, and compatibility. Idioms that are harmless on conventional systems may become problematic once compartmentalization introduces stronger protection semantics. To date, there is no systematic overview of software idioms that become problematic under different compartmentalization semantics.

To fill this gap, we plan to collect problematic idioms and patterns, with their corresponding compartmentalization semantics, and create a taxonomy. This poster presents our agenda for systematically identifying migration-relevant software idioms and their interactions with capability systems. We aim to build a bidirectional taxonomy linking recurring idioms, root causes, and interacting capability primitives. By synthesizing evidence from publications, repositories, and forums, the study will inform future compartmentalization systems and migration tools.

Index Terms—Compartmentalization, capability systems, software security

I. PROBLEM AND MOTIVATION

Software vulnerabilities remain a persistent, evolving problem, and disclosure processes can create periods in which defenders lag behind attackers [1]. Least-privilege design therefore remains a key strategy to reducing compromise impact. Software compartmentalization implements this principle by decomposing systems into isolated components, each granted only required privileges [2]. Instead of giving all code the same ambient authority, it introduces boundaries so that one compromised compartment need not compromise the whole system [3]–[5].

Although used in practice [6], [7], compartmentalization is still not an established engineering practice. Most existing systems were not built with compartmentalization—and even less object capabilities—in mind, while research systems that support it often require substantial manual effort to adapt legacy software [3]. Adapting existing software to compartmentalized designs forces developers to balance correctness, security, complexity, and performance. Existing tools are mostly semi-automated, rely on source-code annotations

or high-level policy input, and still struggle with difficult conversions and real-world edge cases [3], [5].

Prior work shows that part of the difficulty in migrating existing software to compartmentalized systems stems from software idioms that are harmless under conventional architectures but become problematic under stronger protection semantics. Examples include pointer-related assumptions [8], [9], large monolithic functions [10], nested memory allocations [11], and function pointers or nested call structures that complicate partitioning [12].

We target object capability systems because they support compartmentalization by replacing ambient authority with explicit, delegable rights and fine-grained protection domains [9], [13], [14]. The difficulty is that capability systems represent and enforce capabilities differently. For example, In CHERI, the pointer value is itself a capability carrying protection metadata; in RV-CURE, the pointer remains a conventional machine word, while protection metadata is stored separately and reached through a tag in unused pointer bits. Thus, the same software idiom can have different consequences under the two models. CHERI capabilities are wider than ordinary integer pointers and include bounds, permissions, and a validity tag, so code assuming pointers fit conventional integer types may need changes (Listing 1). RV-CURE uses unused pointer bits and therefore avoids this pointer-size mismatch [15].

```
#include <stdint.h>
uint64_t raw = (uint64_t)p; // - Legacy assumes
void *q = (void *)raw; // ptr fits in 64 bits -
uintptr_t raw = (uintptr_t)p; // - CHERI preserve
void *q = (void *)raw; // ptr provenance -
```

Listing 1: Pointer-to-integer round trip in legacy C and CHERI-aware C.

What is still missing is a systematic account of migration-relevant software idioms, their underlying root causes, and the capability primitives with which they interact. Existing empirical studies and porting reports provide important evidence for this gap, but they usually describe blockers within a specific system, platform, or case study [16]. In this work, such reports are treated as input evidence. The contribution is a cross-system synthesis that maps recurring idioms to their root causes and to the capability-system primitives that make them difficult to migrate. This gap motivates the following

research questions:

RQ1. Which idioms across software stacks recur as obstacles during migration to capability systems?

RQ2. How do capability-system primitives shape these obstacles, and in which software-capability-system combinations is automation realistic?

The expected outcome of this systematization is a taxonomy that connects the characteristics of capability systems with the techniques and primitives used to realize them, and with the software idioms that become difficult to migrate under those design choices. The first part of the taxonomy will provide a cross-matching table between capability-system characteristics, their enabling primitives or mechanisms, and the recurring difficult idiom patterns observed in the collected corpus. Where space allows, representative examples will be included for the most prominent cases. The second part of the taxonomy will group the identified idioms by similarity and compare their migration difficulty using a preliminary set of criteria. These criteria are intended as a working definition and may evolve as the taxonomy is refined.

On one hand, this taxonomy can help characterize which classes of software are likely to be difficult to migrate, which idioms recur across software stacks, and where automation is realistic. On the other hand, it can help characterize capability primitives themselves in terms of the migration obstacles, incompatibilities, and adaptation burdens. In turn, this can guide both the design of future compartmentalization systems and auto-compartmentalization tools.

II. APPROACH

Our planned work has three stages: corpus construction, evidence extraction, and creating the taxonomy.

a) Corpus construction: We construct an evidence corpus from three source types. The first is academic publications and technical reports, which describe capability-system designs, compartmentalization mechanisms, and known migration limitations, and may therefore mention problematic idioms. For example, Gudka et al. [5] report that “The `dlopen()` call will fail when in the sandbox,” illustrating how lazy initialization and dynamic library loading can become compartmentalization obstacles. For this source type, we handpicked 36 papers and technical reports as a seed covering different capability-system characteristics.

We do not start from a fixed list of idioms or restrict the study to a single application family. Instead, idioms are treated as open-ended findings to be discovered during evidence extraction. Our initial seed set is therefore selected to cover a diverse range of capability-system characteristics, so that the resulting taxonomy can capture how different protection primitives and semantics expose different migration obstacles. Starting from the seed papers, we applied two rounds of backward and forward snowballing [17] using OpenAlex [18] to retrieve publication metadata and citation graphs, yielding 12,621 papers. Keyword filtering removed clearly irrelevant papers and left 7,525 papers. We will next apply context-aware

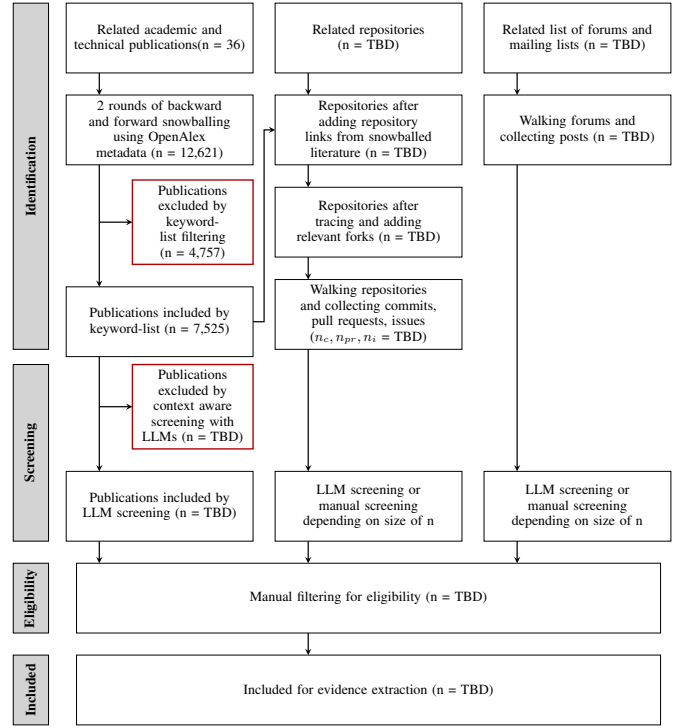


Fig. 1. PRISMA Flow Diagram for our Corpus Construction (§II-0a) and Evidence Extraction (§II-0b)

LLM screening to titles and abstracts, then manually decide whether the remaining resources contain relevant idioms.

Challenge 1: Evidence Identification

Relevant idioms may appear in many parts of a publication, from evaluation to future work, and not necessarily as focal point of the publication.

The second source type consists of open-source repositories, their commit histories, and issue trackers of software ported to capability systems. These provide concrete examples of porting problems. For example, one CHERI BSD issue [19] points to an `INT2PTR` cast, illustrating how repository discussions can expose migration failures such as pointer-integer conversions that violate CHERI’s capability semantics [20]. Starting from the repositories of capability systems, we will collect pull requests, issues, and commit messages. Depending on the corpus size, we will apply LLM screening to decide whether the collected resources are relevant. We will then apply manual filtering to determine eligibility.

Challenge 2: Repository Selection and False Positives

Not all capability-system repositories, provide useful migration evidence; screening must distinguish idiom constraints from ordinary engineering noise.

The third source type consists of forums, mailing lists, and other developer-discussion platforms, which can reveal practical problems and informal workarounds absent from papers and documentation. For example, one FreeBSD mailing-list post [21] reports: “Do not try to build `cloudabi32` for `pc98`. Should unbreak `tinderbox`.” Even brief build reports

