# Code-centered kernel compartmentalization in CheriBSD

Konrad Witaszczyk, Department of Computer Science and Technology, University of Cambridge, March 2025

The CHERI-extended CheriBSD kernel based on the FreeBSD kernel, as other monolithic kernels, consists of millions of trusted lines of code that are compiled into the kernel binary or separate kernel modules. The kernel binary itself includes code from over 1,700 ELF object files linked together. In this research, I explore compartmentalization of the CheriBSD kernel that focuses on its code (Figure 1) and aims to split the kernel code into compartments that can call functions of other compartments only if a system-defined policy allows such a call.

/boot/kernel/kernel /boot/kernel/zlib.ko /boot/kernel/kernel /boot/kernel/zlib.ko /boot/kernel/kernel /boot/kernel/zlib.ko compartment load/unload compartment load/unload exception handler k. c. allocator allocator routines vector routines: executive kernel executive k. c. kernel exception handler exception handler zlib exception handler zlib mod.c k. c. allocator kernel module vector vector per-CPU context per-CPU context per-CPU context kernel compression allocator: kernel handler handler k. c. handler algorithms k. c. crypto zcalloc.c ... ... ... compressor compression compartment k. c. exception handler exception handler algorithms: crypto deflate.c k. c. crypto k. c. compressor (a) 0 compartments: trusted kernel; inflate.c kernel k. c. compressor trusted zlib and other kernel modules (\*.ko). k. c. (b) 5 compartments: trusted executive; untrusted rest of kernel; 3 untrusted zlib kernel module. (c) compartments optimised for security and performance: trusted executive; separated untrusted kernel subsystems trusted component and modules. code-centered kernel compartments untrusted component



Figure 1. Incremental code-centered compartmentalization approach in a monolithic kernel. Case b presents the current state.

### Static and dynamic linking for compartmentalization

The kernel and kernel modules are statically linked into binaries that consist of object files grouped into compartments. When relocating their symbols at run time, the dynamic kernel linker wraps the symbols with trampolines that implement compartment switching.

## **Calls between compartments**

The trampolines (Figure 2) use architectural features (e.g., the Executive and Restricted modes in Arm Morello) to protect capability registers from being manipulated by an untrusted compartment. Those features are specific to a CHERI-extended architecture and are co-designed with software.

## **Optimal compartment boundaries**

The security and performance characteristics of the system can significantly differ depending on the choice of the Trusted Computing Base and untrusted compartment boundaries. The search for the optimal compartmentalization policy requires extensive experiments and engineering.



Figure 2. An example function call between compartments using a trampoline. The blrr instruction branches into the Restricted mode using a capability without the Executive permission bit set making the csp register refer to rcsp\_el0 instead of csp\_el1



**Hypothesis:** Hardware-assisted compartmentalization with CHERI can significantly reduce the risk of vulnerabilities being exploited in a monolithic kernel with minor implementation and performance costs.

#### **Contact** Konrad.Witaszczyk@cl.cam.ac.uk



Approved for public release; distribution is unlimited. Sponsored by the Defense Advanced Research Projects Agency (DARPA) and the Air Force Research Laboratory (AFRL), under contract FA8750-24-C-B047 ("DEC") as part of the I2O CPM research program. The views, opinions, and/or findings contained in this report are those of the authors and should not be interpreted as representing the official views or policies of the Department of Defense or the U.S. Government.