

Winning the war in memory using CHERI capabilities


Robert N. M. Watson, **Simon W. Moore**, Peter G. Neumann, Peter Sewell

Hesham Almatary, Jonathan Anderson, John Baldwin, Hadrien Barrel, Ruslan Bukin, David Chisnall, Nirav Dave, Brooks Davis, Lawrence Esswood, Nathaniel W. Filardo, Khilan Gudka, Alexandre Joannou, Robert Kovacsics, Ben Laurie, A. Theo Marketos, J. Edward Maste, Alfredo Mazzinghi, Alan Mujumdar, Prashanth Mundkur, Steven J. Murdoch, Edward Napierala, Kyndylan Nienhuis, Robert Norton-Wright, Philip Paeps, Lucian Paul-Trifu, Alex Richardson, Michael Roe, Colin Rothwell, Hassen Saidi, Peter Sewell, Stacey Son, Domagoj Stolfa, Andrew Turner, Munraj Vadera, Jonathan Woodruff, Hongyan Xia, and Bjoern A. Zeeb

University of Cambridge and SRI International
SILM Summer School, Inria, Rennes – 11 July 2019

Motivation – The Eternal War in Memory*

- Many security vulnerabilities exploit memory safety violations



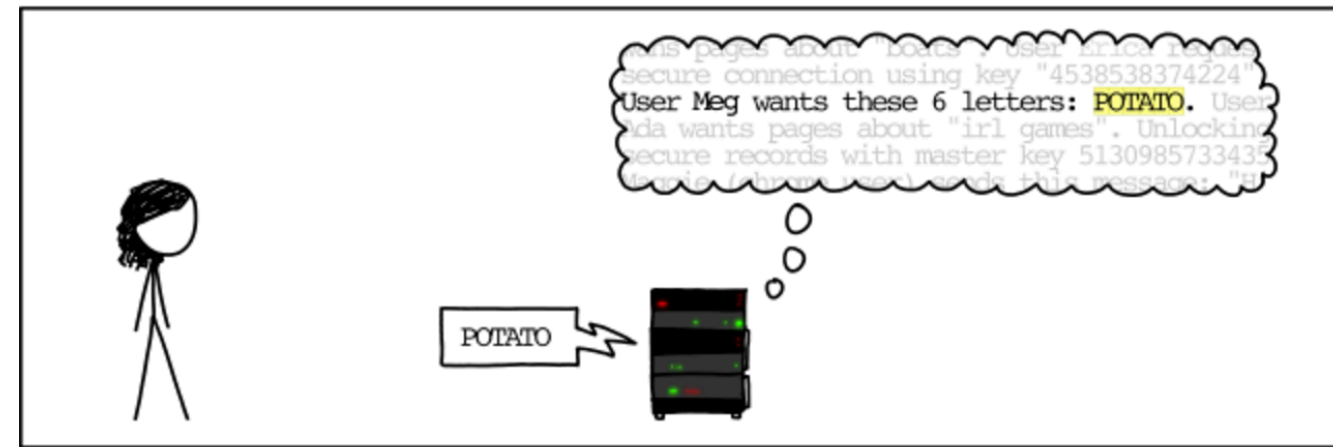

*Title based on Oakland 2013 paper: SoK: Eternal War in Memory, László Szekeres, Mathias Payer, Tao Wei, Dawn Song

HOW THE HEARTBLEED BUG WORKS:

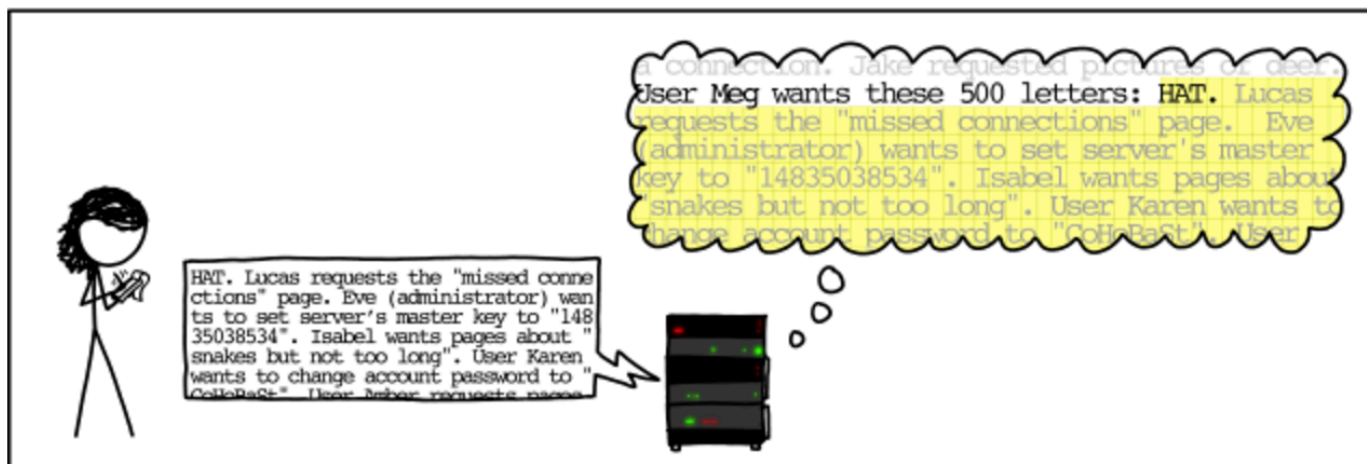
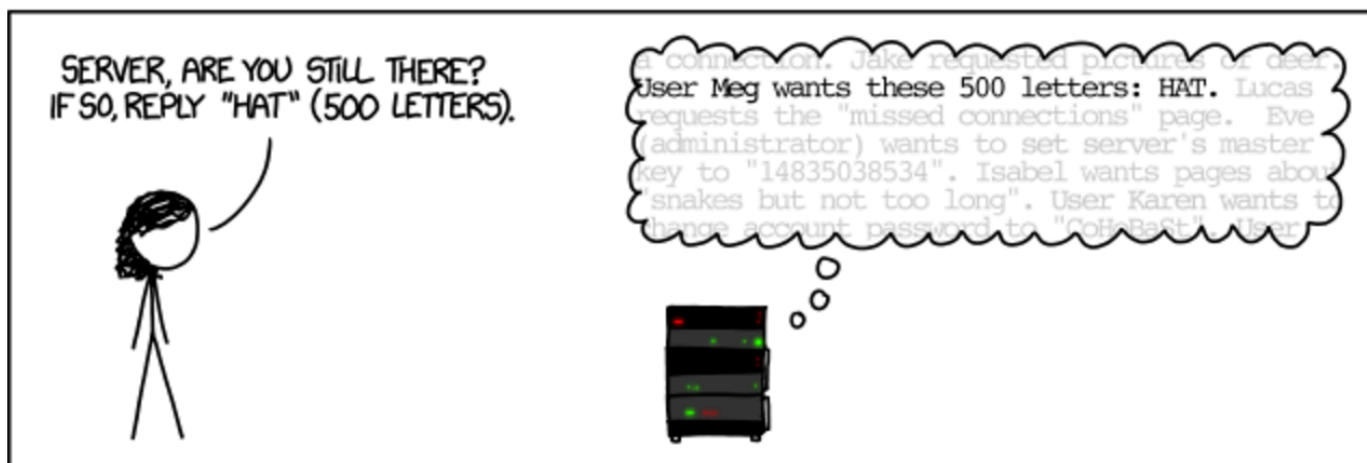
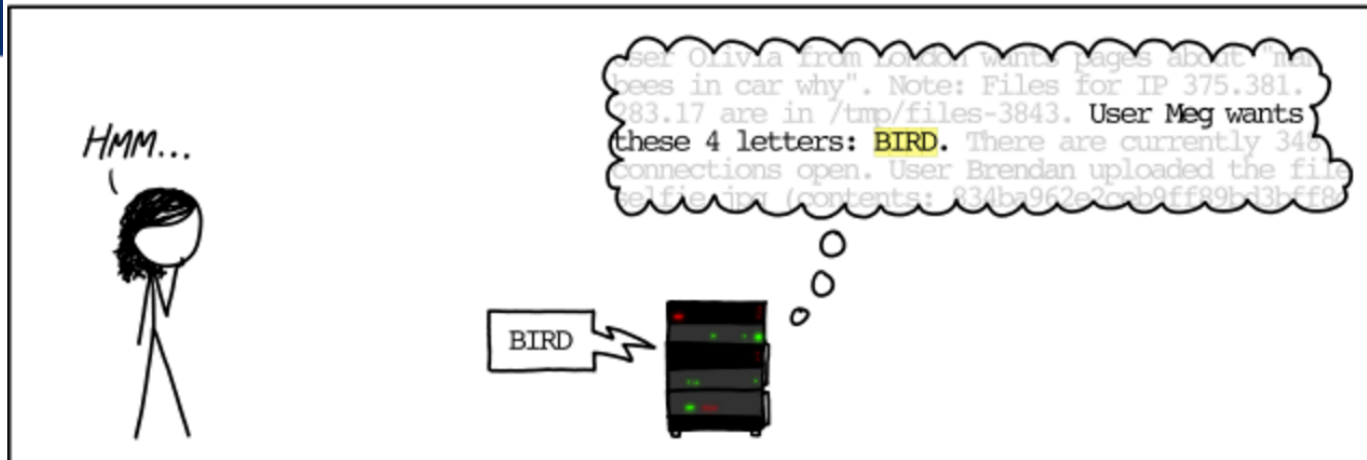
Example 1



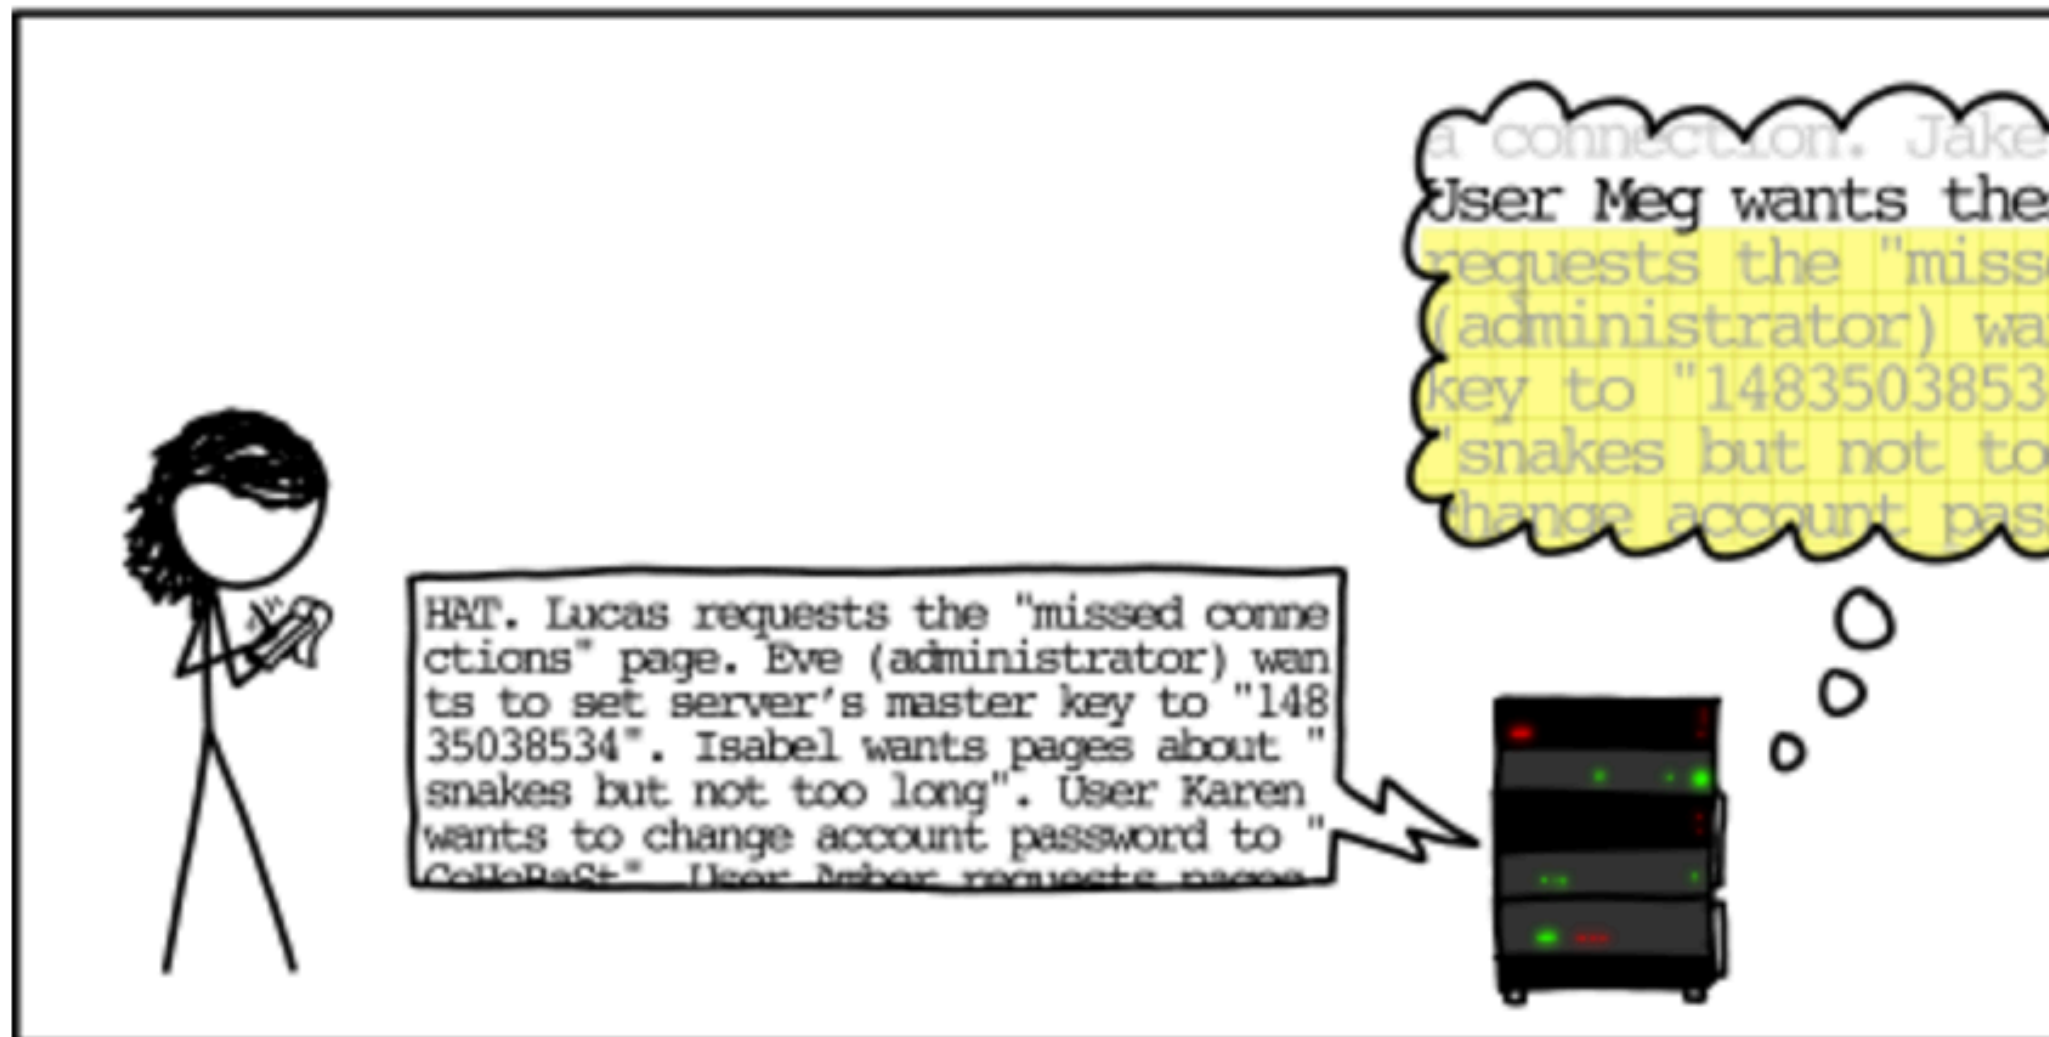
HeartBleed



source: <http://xkcd.com/1354/>



source: <http://xkcd.com/1354/>



Went wrong? How do we do better?

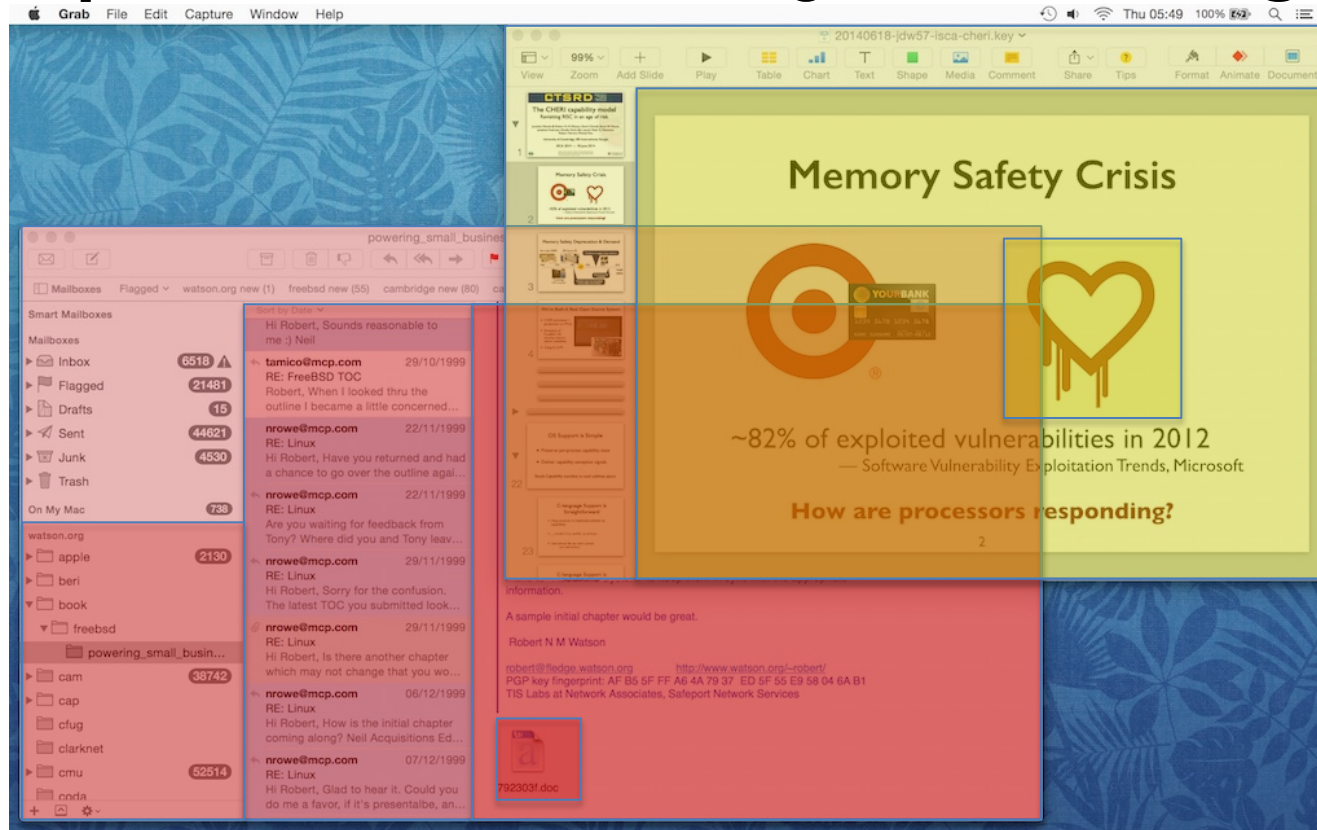
- Classical answer:
 - The programmer forgot to check the bounds of the data structure being read
 - Fix the vulnerability in hindsight – one line fix:
`if (l+2+payload+l6 > s->s3->rrec.length) return 0;`
- Our answer:
 - Preserve bounds information during compilation
 - Use hardware (CHERI processor) to dynamically check bounds with little overhead and guarantee pointer integrity & provenance

Example 2: how to reduce the attack surface?

- The software attack surface keeps getting bigger
 - Applications just keep getting larger
 - Huge libraries of code aid rapid program development
 - Everything is network connected
- This aids the attacker: an expanding number of ways to break in

Application-level least privilege

Software compartmentalization decomposes software into **isolated compartments** that are delegated **limited rights**



Able to mitigate not only **unknown vulnerabilities**, but also **as-yet undiscovered classes of vulnerabilities and exploits**

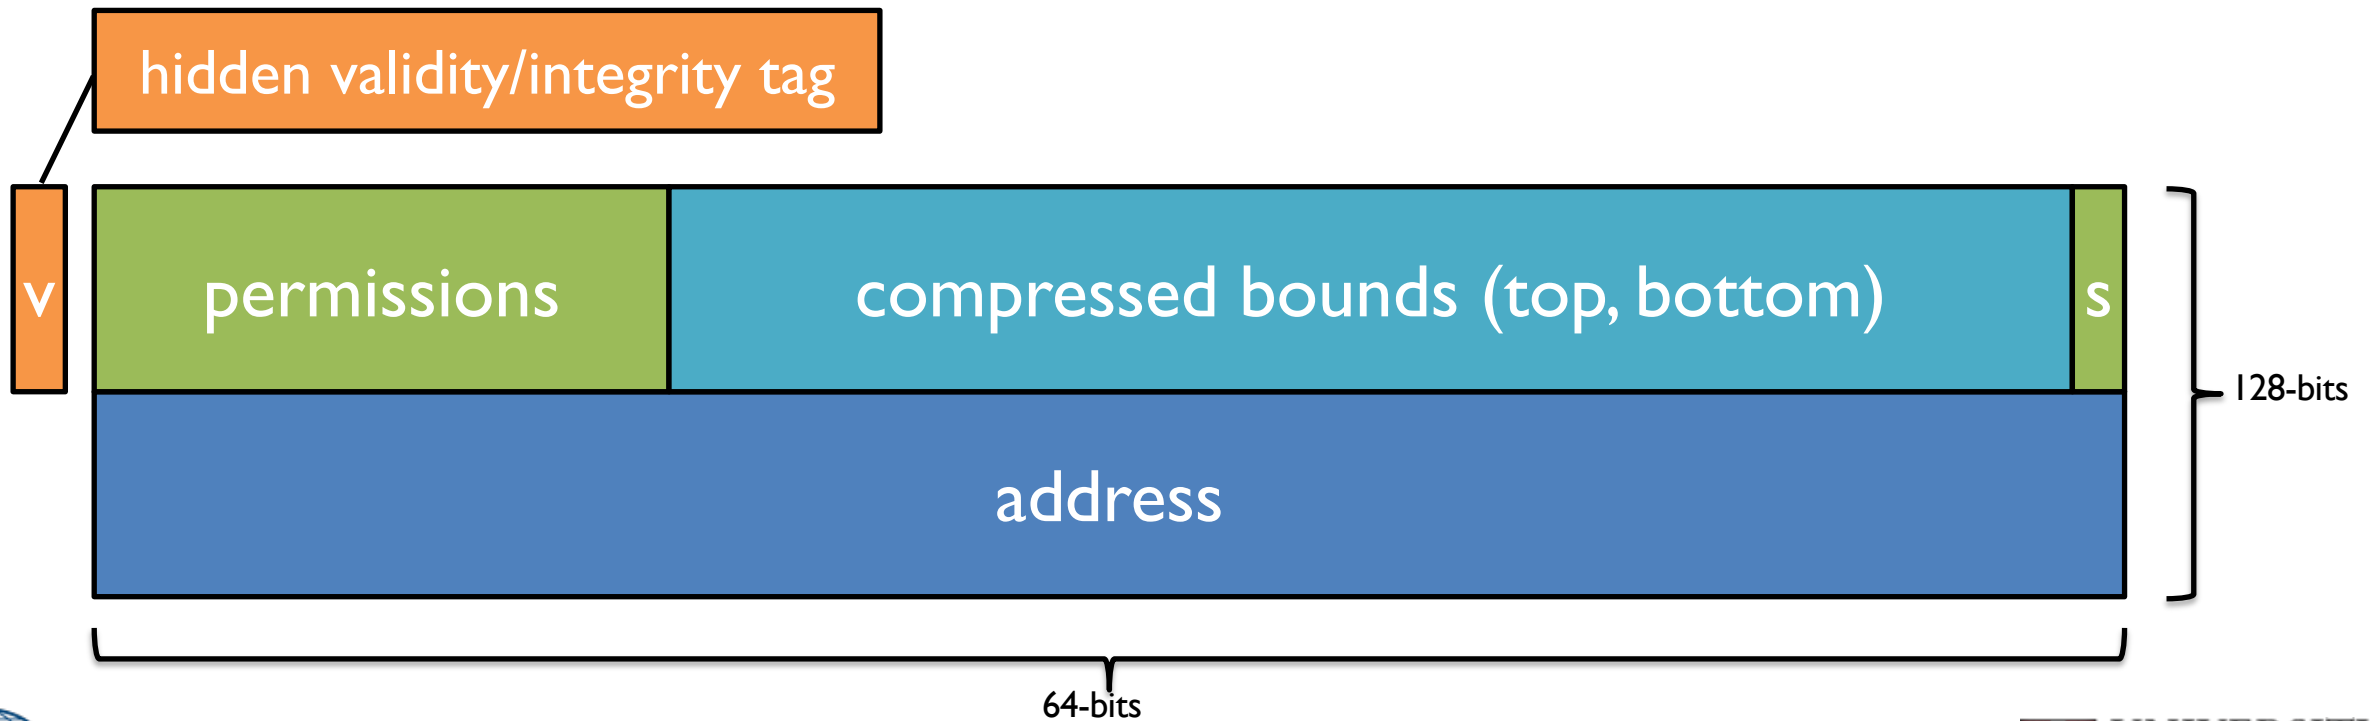
Principles CHERI helps to uphold

- The **principle of intentional use**
 - Ensure that software runs the way the programmer intended, not the way the attacker tricked it
 - Approach: guaranteed pointer integrity & provenance, with efficient dynamic bounds checking
- The **principle of least privilege**
 - Reduce the attack surface using software compartmentalization
 - Mitigates known and unknown exploits
 - Approach: highly scalable and efficient compartmentalization

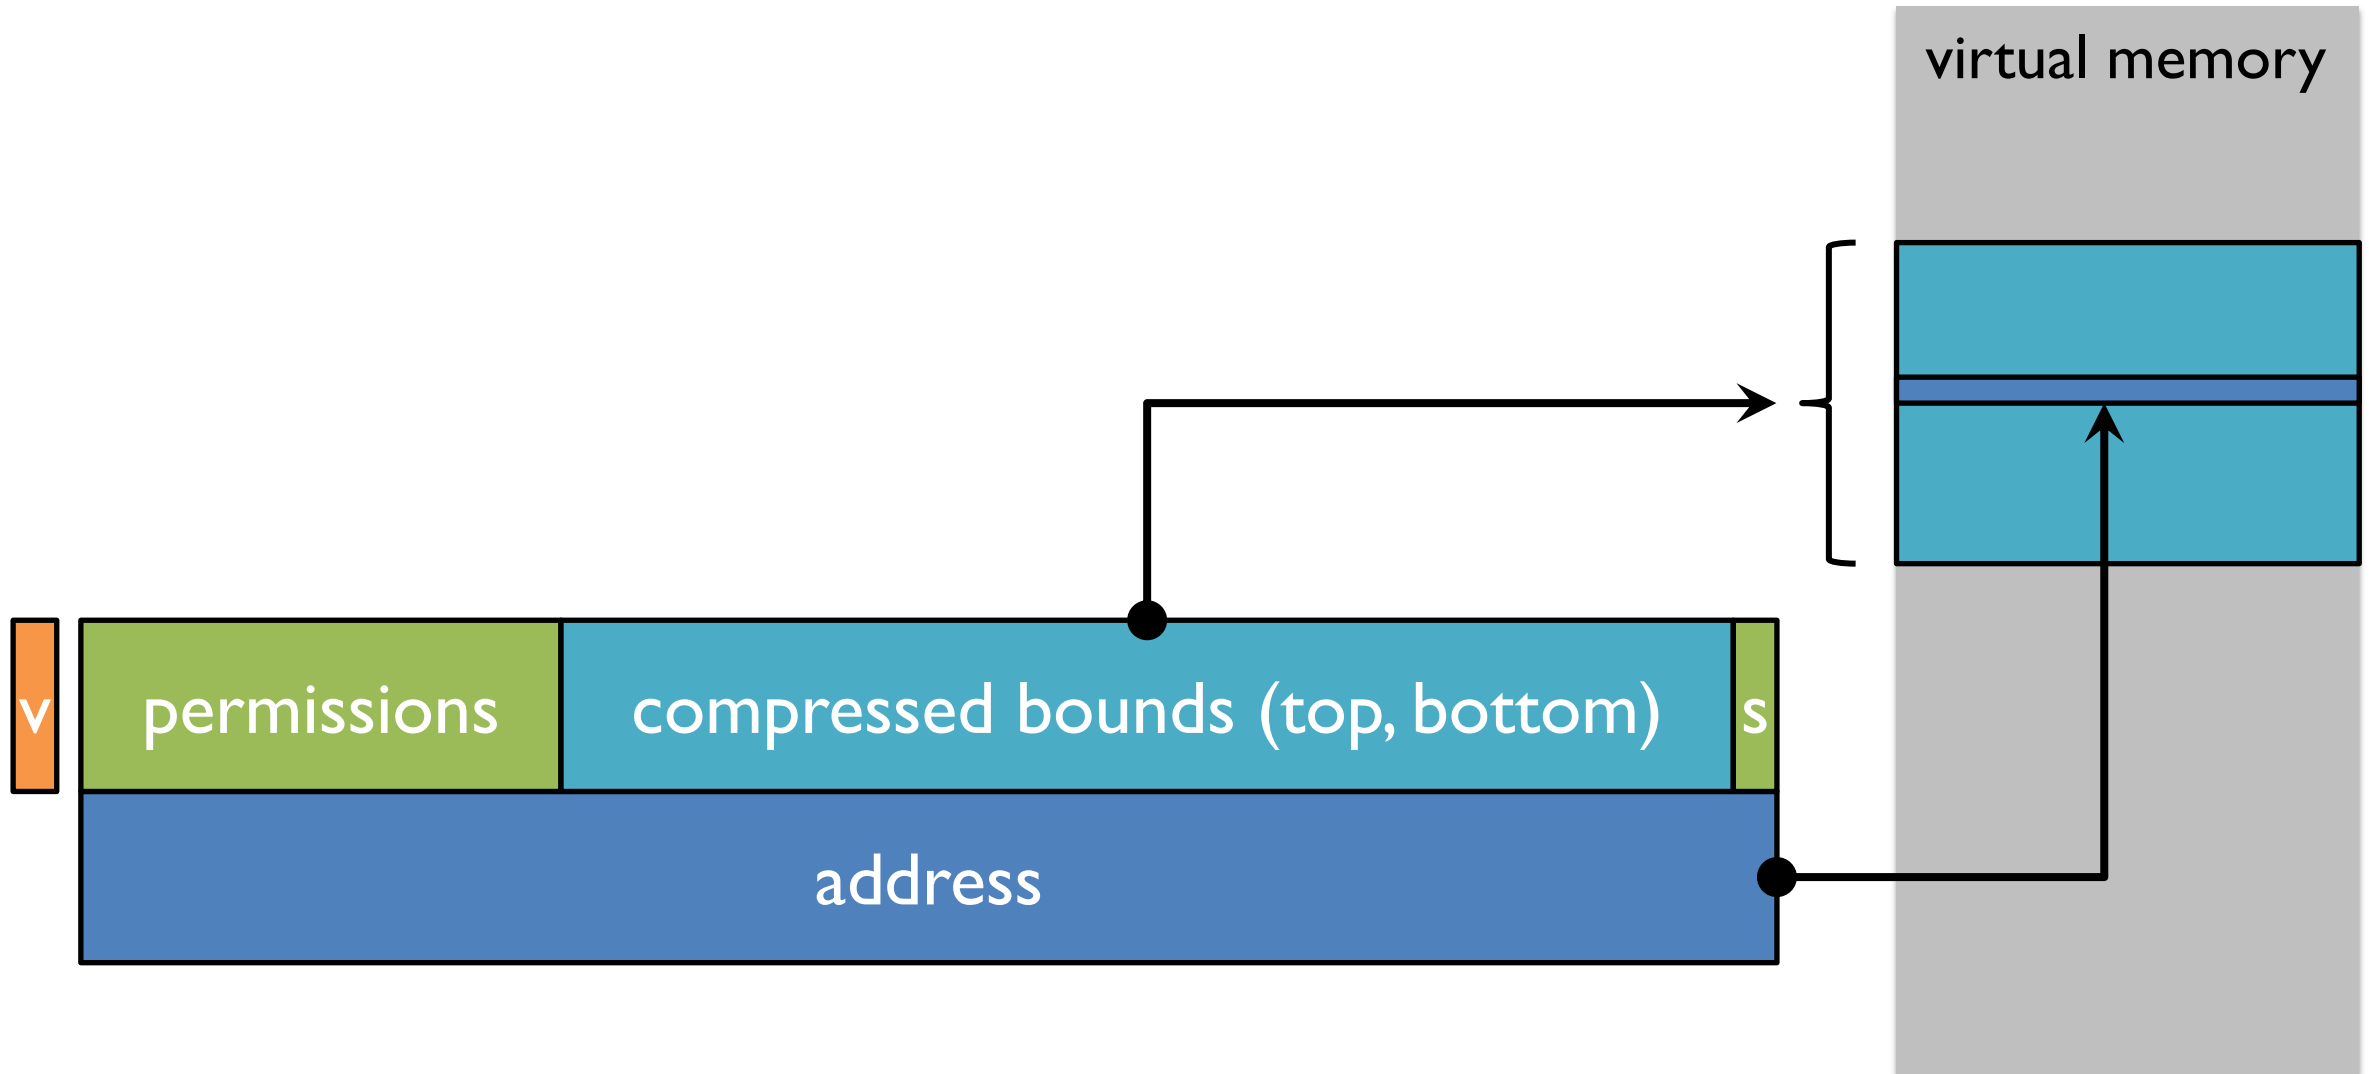
CHERI HARDWARE ARCHITECTURE

A new type – the **Capability**

- CHERI Capability = bounds checked pointer with integrity
- Held in memory and in (new or extended) registers



A new type – the **Capability**



New Instructions

- Memory access
 - Loads and stores via a bounds checked capability
 - Exception if address is out of range
- Guarded manipulation of capabilities

- Decrease bounds

- Decrease permissions

- Adjust the address

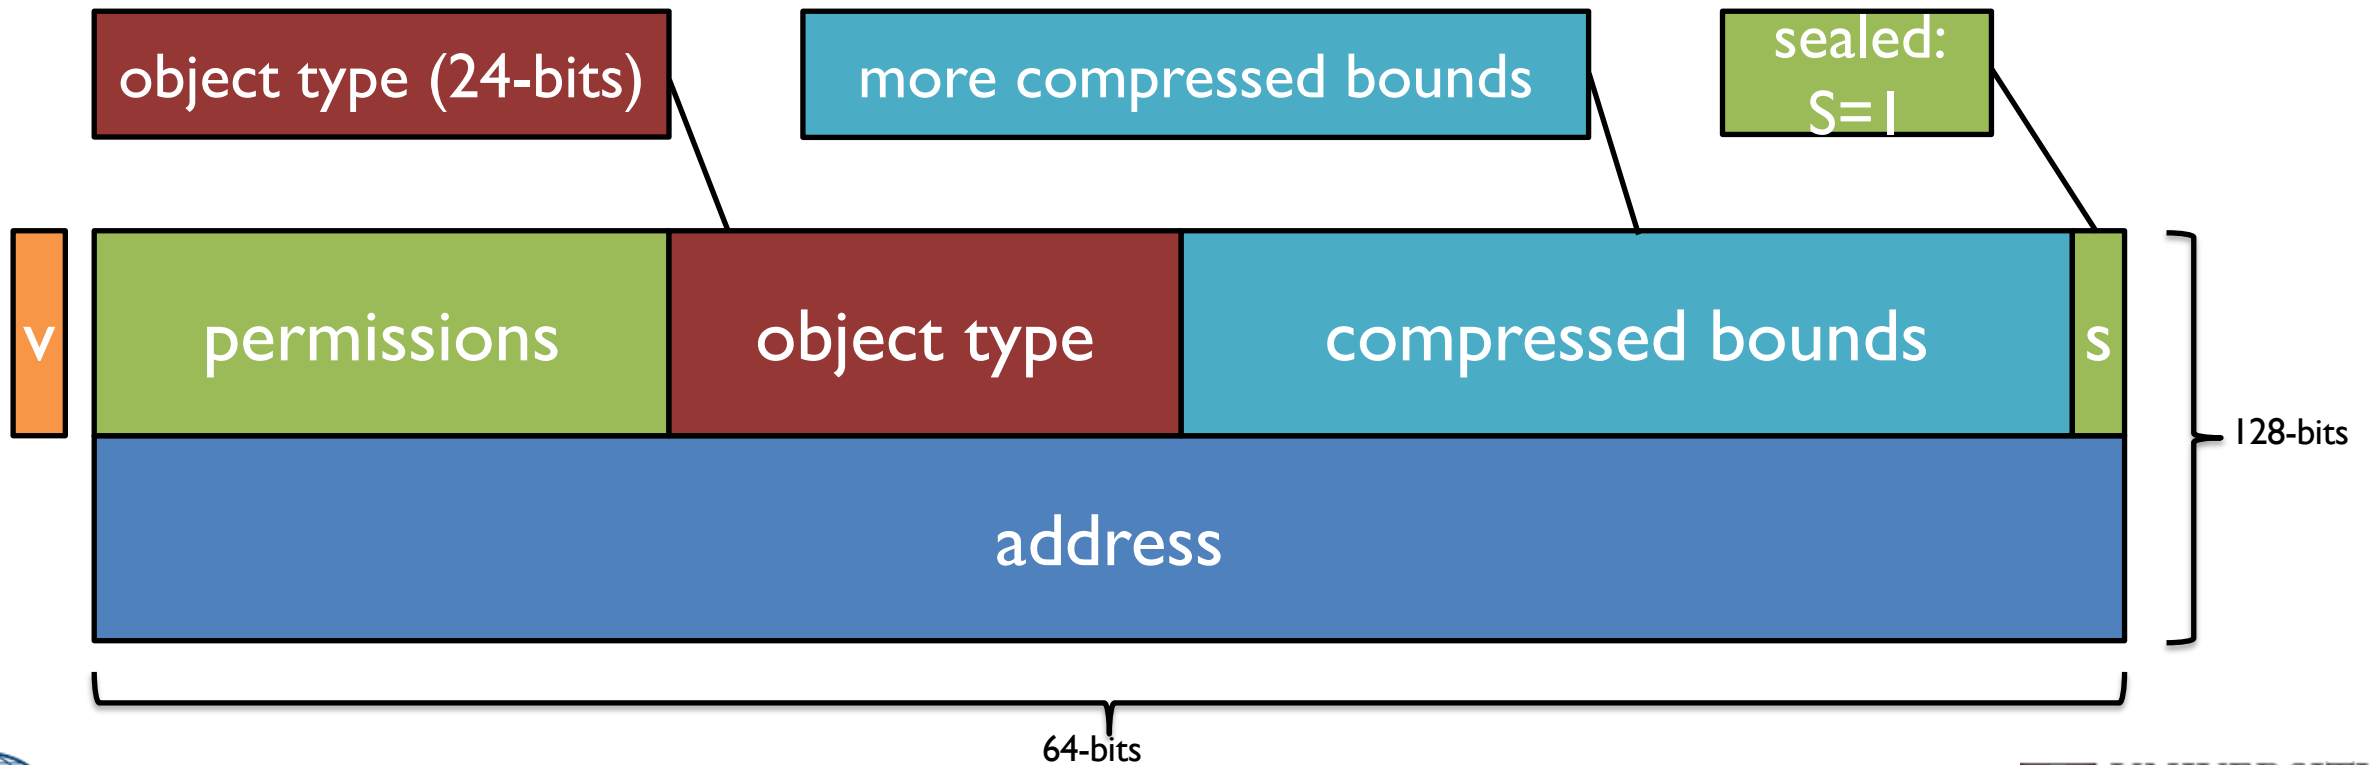
- Extract/test fields

monotonic decrease in rights guaranteed
by formally verified hardware

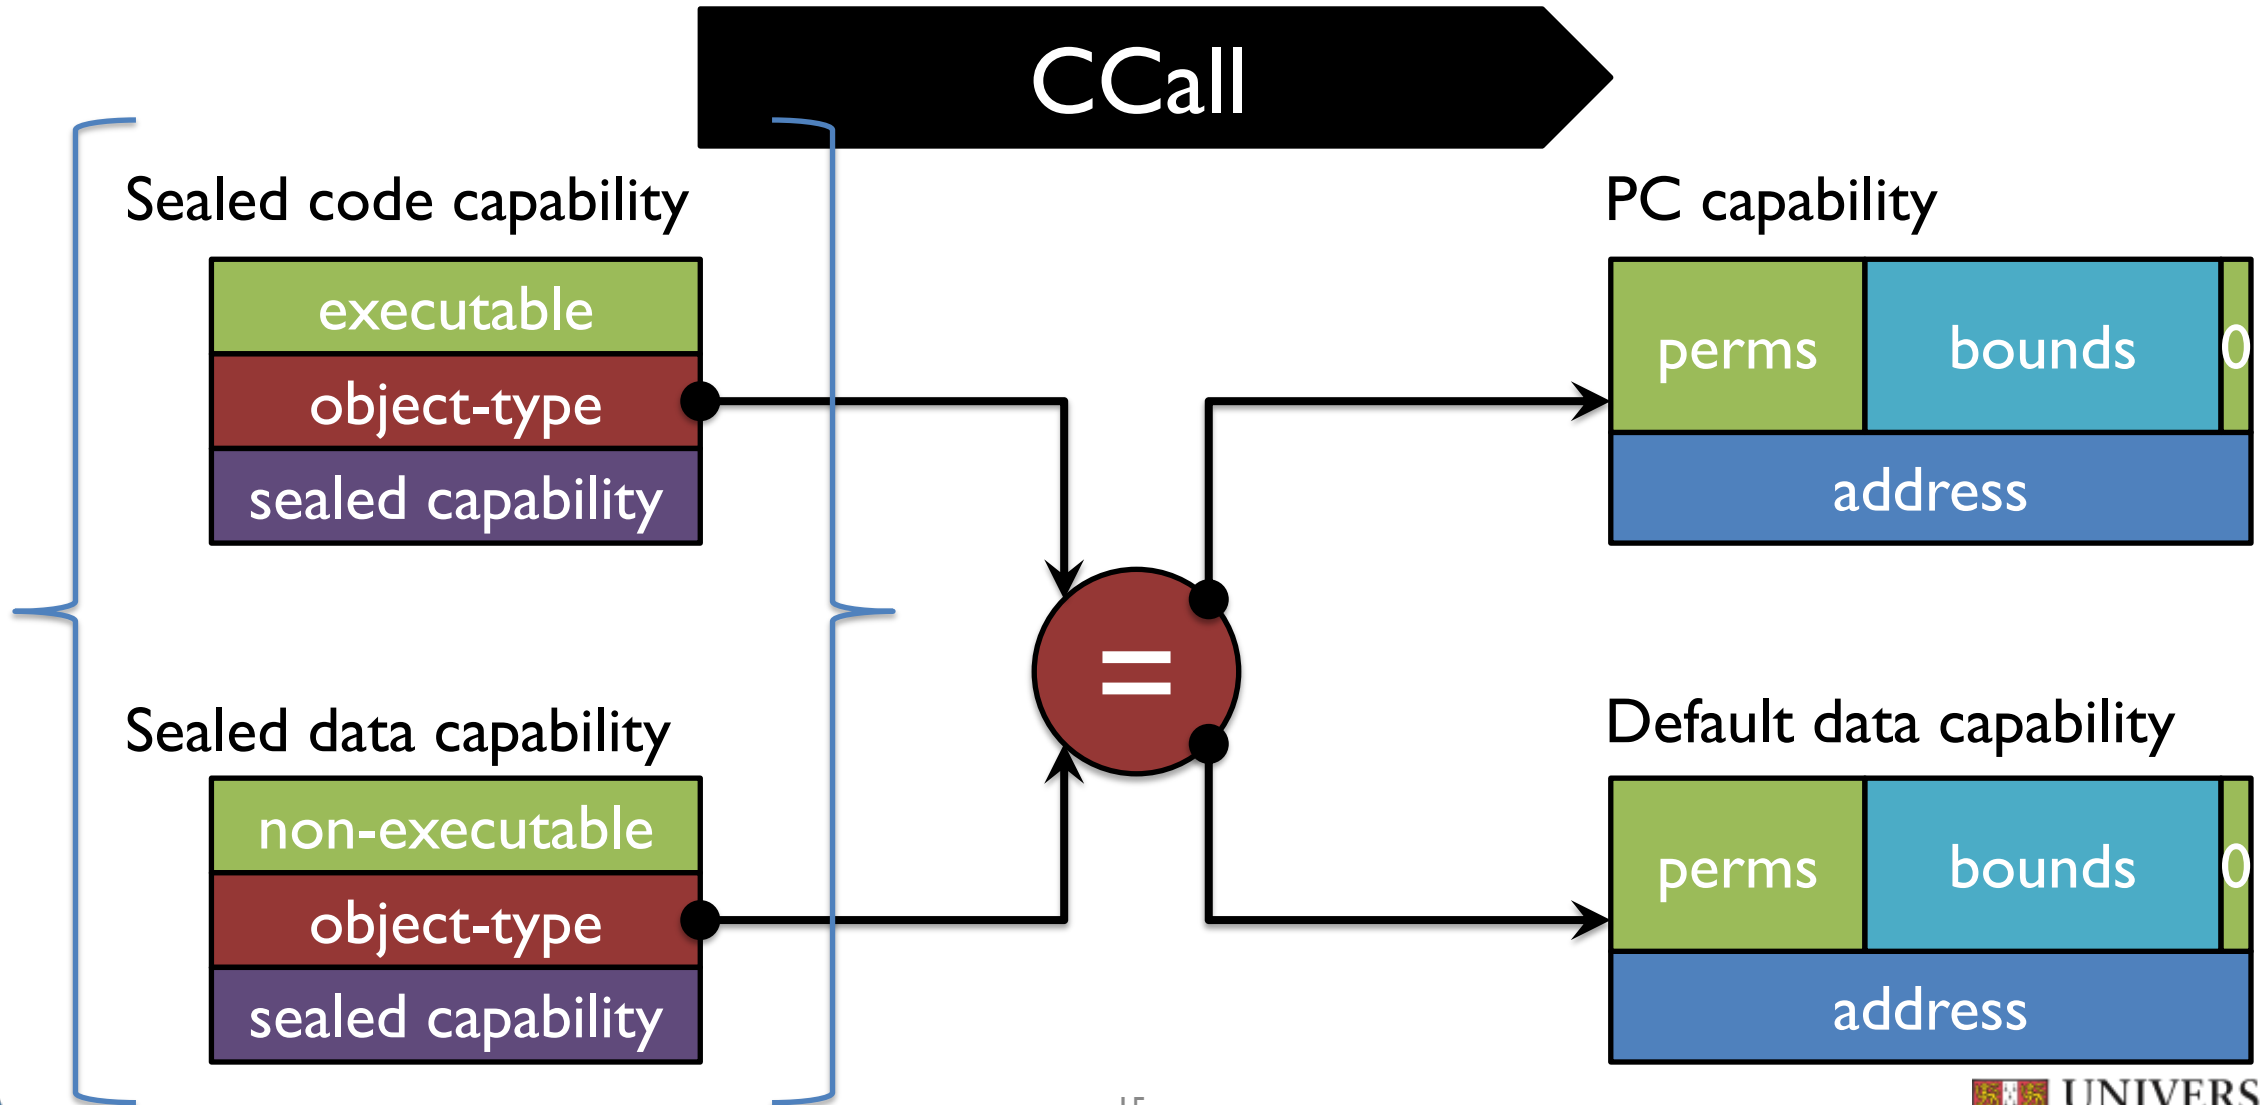
critical property for security

Sealed Capabilities for Compartmentalization

- Sealed capabilities are none dereferencable capabilities
- Have to be unsealed (e.g. inside a compartment) before use



Calling a Compartment




SOFTWARE MODELS

Background to CHERI Software Models

- Machine-level capabilities and instructions provide the building blocks on which new abstract capability software models can be built
- Analogy:
 - Machine-level translation lookaside buffer (TLB) and page table walker enables the OS to represent virtual memory
 - Virtual memory can then be used in different ways to impose new security features, e.g. guard pages

Low-level CHERI software models



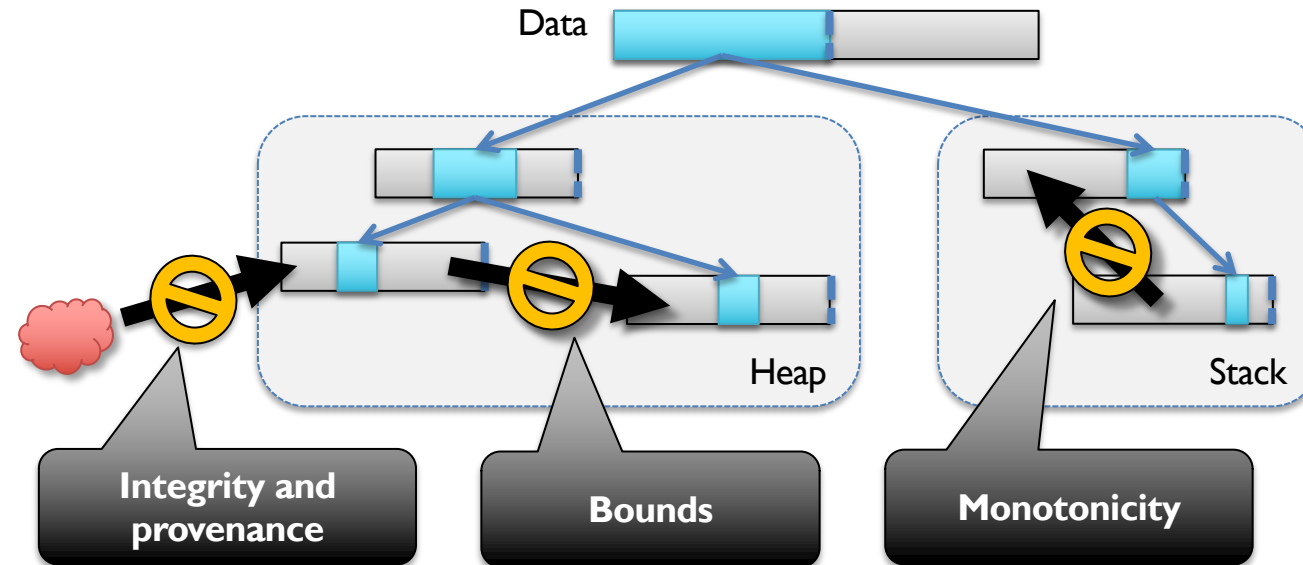
- **Source and binary compatibility: C-language idioms, multiple ABIs**
 - **Unmodified code:** Existing code runs without modification
 - **Hybrid code:** E.g., used in return addresses, for annotated data/code pointers, for specific types, stack pointers, etc.
 - ... But “hybrid” is a spectrum: many different choices for manual and automatic selection of integers vs. capabilities, API and ABI impacts
 - **Pure-capability code:** Ubiquitous data- and data-pointer protection. Not interoperable with legacy code due to changed pointer size.
- **CHERI Clang/LLVM compiler prototype** generates code for all three

Pure Capability Code → Needs CheriABI

- CheriABI
 - Compatibility layer to the OS
 - Allows capabilities to be used in place of pointers
 - A bit like a 32-bit compatibility layer for a 64-bit OS
- Result – we can now recompile large corpuses of C code into a pure capability form with virtually no code changes
- Award winning paper at ASPLOS 2019:
CheriABI: Enforcing Valid Pointer Provenance and Minimizing Pointer Privilege in the POSIX C Run-time Environment

Capabilities for Bounds Checking and Integrity

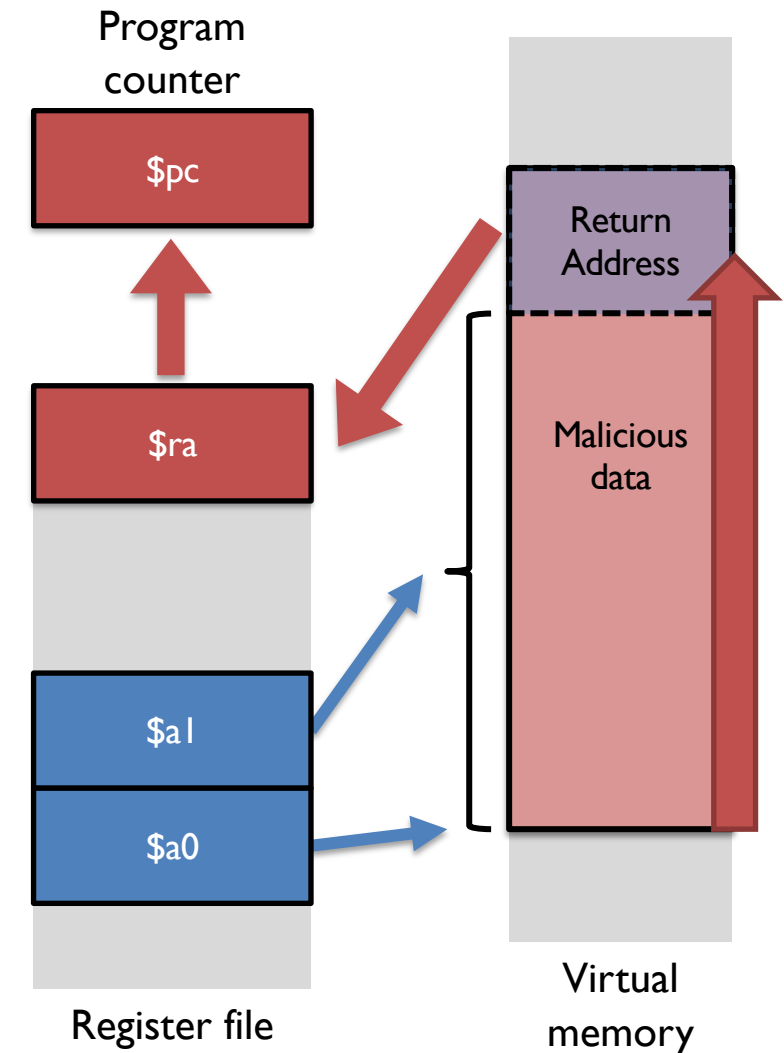
- Pure capability code – all pointers become capabilities
- Compiler + malloc() derive bounds for objects
- Strong pointer provenance and integrity properties (validity tag)




- Mitigates buffer overflow/overread vulnerabilities with no code change!

Capabilities for Control-Flow Robustness

- Capabilities used for return addresses
- Integrity bit mitigates code reuse attacks:
 - ROP – Return Oriented Programming
 - JOP – Jump Oriented Programming
- Much better than current statistical technique ASLR (Address Space Layout Randomisation)



Summary of Capability Protections



Valid userspace pointer set – pointers not generated using derivation rules are not part of the valid provenance tree and will not be dereferenceable

Pointer privilege reduction – capabilities allow pointers to carry specific privileges, which can be minimized with OS, compiler, and linker support

Foundation for higher-level models such as **software compartmentalization**

Compartmentalisation


- Compartment can be described using a sealed pair of capabilities: (program counter, default data capability)
- CCall providing the domain transition
- Allows a number of abstract software models:
 - Library compartmentalisation, e.g. of risky or legacy (non-cap.) code
 - Process-based compartmentalisation within an application can be replaced by much more efficient capability-based protection
 - Same virtual address space (more efficient TLB usage)
 - Very similar software model (easy to port code)

RESULTS

First we made it work – Demo tablet platform



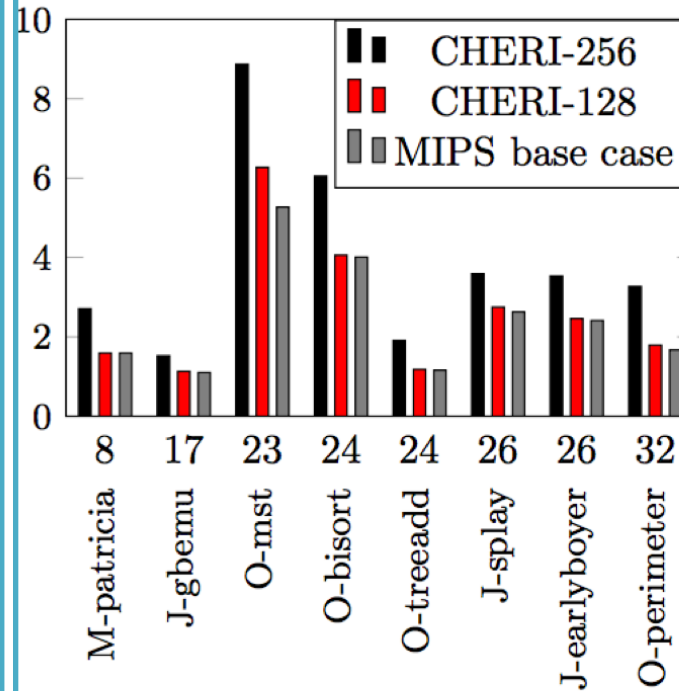
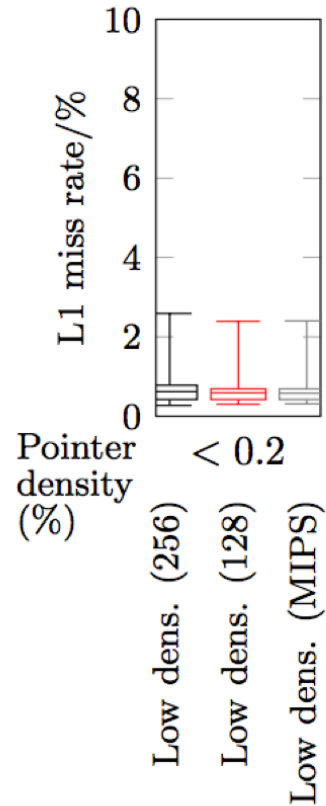
Red Team Evaluation by MIT Lincoln Labs



CHERI mitigates
Heartbleed exploit!

Memory-protection performance

Collection of low pointer-density benchmarks from MiBench



High pointer-density benchmarks

(M) MiBench

(O) Olden

(J) Octane JavaScript

L1 cache miss rate for CHERI 256, CHERI-128, and MIPS

CheriABI: A full pure-capability OS userspace

- Complete memory- and pointer-safe FreeBSD C/C++ userspace
 - **System libraries:** crt/csu, libc, zlib, libxml, libssl, ...
 - **System tools and daemons:** echo, sh, ls, openssl, ssh, sshd, ...
 - **Applications:** PostgreSQL, nginx; bringing up WebKit (C++)
- **Valid provenance, minimized privilege for pointers, implied VAs**
 - Userspace capabilities originate in **kernel-provided roots**
 - Compiler, allocators, run-time linker, etc., **refine** bounds and perms
- Trading off **privilege minimization, monotonicity, API conformance**
 - Typically in memory management – realloc(), mmap() + mprotect()

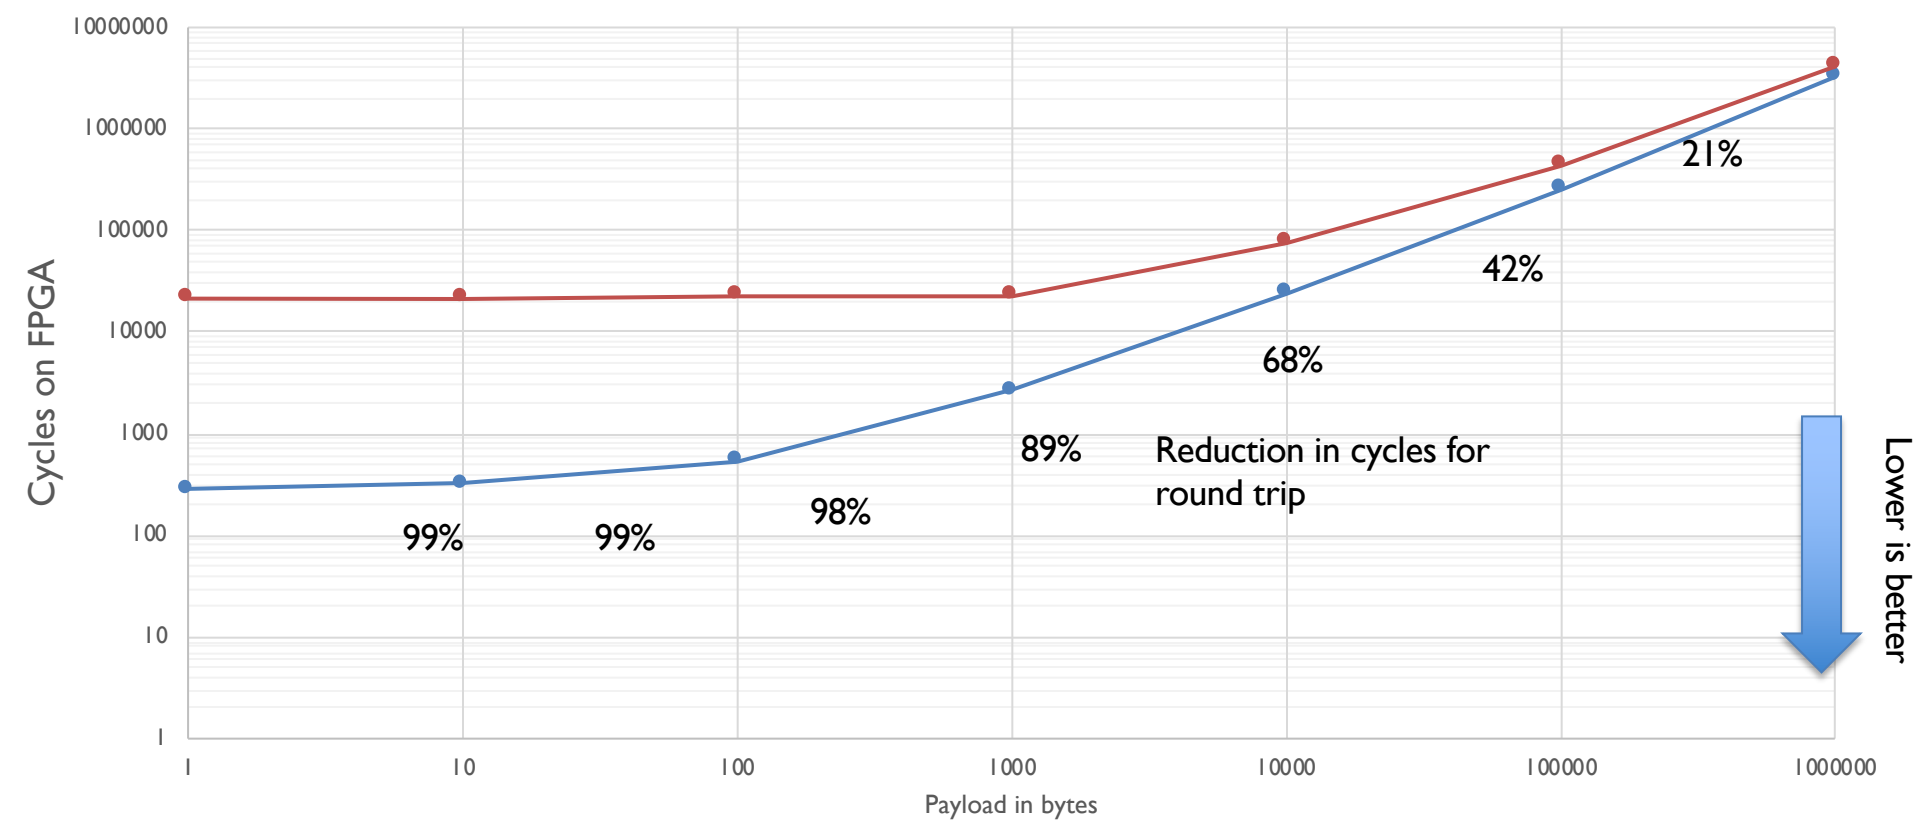
Evaluating memory-protection compatibility

- Prototyping approach:
 - “pure-capability” **C compiler** (Clang/LLVM)
 - **full OS** (FreeBSD) that use capabilities for all explicit or implied userspace pointers
- Observations:
 - **Little or no software modification** (BSD base system + utilities)
 - Small changes to source files for 34 of 824 programs, 28 of 130 libraries
 - Overall: modified ~200 of ~20,000 user-space C files/header

CHERI vs. Process-based Compartmentalization

(Early IPC ping-pong microbenchmark results)

Co-process vs. pipe(2) ping-ping
 Memory-copy semantics with multi-byte payload



The fine print: Cycles include IPC setup, amortised over 10,000 iterations of the IPC loop. Both processes use the pure-capability ABI using 256-bit capabilities. 272-entry TLB, 32K L1 I-Cache, 32K L1 D-Cache, 256K L2 Cache.

—●— Co-process —●— pipe(2)
 30



CURRENT RESEARCH DIRECTIONS

Generalising CHERI support for many ISAs

- 64-bit MIPS for pragmatic reason: needed a 64-bit RISC ISA in late 2010
- Generic CHERI support doesn't mean that all implementations need to be identical
 - E.g. **portable virtual-memory semantics** and **UNIX process model** despite (quite) different MMUs across architectures
- **Architectural abstraction**: Lift CHERI properties above ISA
- **Architectural localization**: E.g., ISA choices, opcode approaches, exceptions, page tables, ... → **architecture-specific specifications**
- Currently working on **CHERI-ARM** and **CHERI-RISC-V** variants
- Currently exploring interaction with virtualization for servers

Portability implications for software

- **CHERI Clang/LLVM**
 - Modest pointer/capability abstraction improvements in front-end and IR
 - Adapt target back-ends to teach them about capability code generation
 - Optimize for architecture-specific code generation
 - Optimize for available microarchitectures
- **CheriBSD** (CHERI support in FreeBSD)
 - More clear machine-independent / machine-independent split
 - Shift to hybrid capability C in the kernel to improve machine independence
 - Various MD kernel updates: boot code, exceptions, PMAP, ...
 - Clean up APIs, header separation, architecture abstraction
 - Various userspace updates: rtd, libcheri, CRT/CSU, ...

Many other research questions

- Can we efficiently impose CHERI protection mechanisms on I/O devices and accelerators?
 - See Thunderclap work on I/O security (<http://thunderclap.io/>)
- Does CHERI make managed languages (e.g. Rust) safer or faster (e.g. through efficient dynamic checks)?
- Does fine-grained compartmentalisation help mitigate fault injection attacks?
- Can CHERI help to mitigate speculative execution attacks?
- Can CHERI be used for enclaves?

Conclusions

- CHERI provides the hardware with more semantic knowledge of what the programmer intended
 - Toward the **principle of intentional use**
- Efficient **pointer integrity** and **bounds checking**
 - Eliminates buffer overflow/over-read attacks (finally!)
- Provide scalable, efficient compartmentalisation
 - Allows the **principle of least privilege** to be exploited to **mitigate known and unknown attacks**
 - Large performance improvement over process-based compartmentalisation
- **Working with industry to bring the technology to market**
- Thanks to sponsors: DARPA, ARM, Google, EPSRC, HEIF, Isaac Newton Trust, Thales E-Security, HP Labs



<https://www.cl.cam.ac.uk/research/security/ctsrd/>

Additional Topics

1. Our verification and test strategy
2. How to build efficient tagged memory
3. Compressed capabilities
4. How CHERI helps to mitigate speculative execution attacks
5. **THUNDER**⚡**CLAP**: The Perils of Peripherals