



CHERI tablet and rack-mount BlueHive array, based on Terasic's DE4 FPGA board. CHERI CPUs support fine-grained compartmentalization, mitigating broad classes of known and unknown vulnerabilities.

CTSRD is developing a principled, formally-supported, robust, programmerfriendly, high-performance and incrementally adoptable hardware/software platform designed for efficient implementation of the principle of least privilege. Hardware and software security structures and design principles are reinforced by:

- Capability Hardware Enhanced RISC Instructions (CHERI)
- Security Oriented Analysis of Application Programs (SOAAP)
- Smten formal verification suite for Bluespec HDL
- Temporally Enforced Security Logic Assertions (TESLA)

CTSRD adopts a hybrid approach, able to run existing C-language operating systems and application software while supporting gradual adoption of its novel protection features for critical Trusted Computing Bases (TCBs) and high-risk components. CTSRD allows programmers to wipe the slate clean incrementally.



Capsicum and the application compartmentalization motivation

Programmers are turning to application compartmentalization to mitigate inevitable vulnerabilities: software is decomposed into sandboxed components, each with only the rights it requires to function. This approach employs the principle of least privilege: as granularity increases, rights delegated to individual sandboxes decrease. As vulnerabilities are exploited, only the rights of the affected component are leaked, forcing attackers to exploit many more vulnerabilities to accomplish the same goals.

The Capsicum hybrid capability model developed by Cambridge and Google blends contemporary OS design with capability system security, addressing semantic mismatches between OS features and application compartmentalization. Rights are delegated using capabilities, unforgeable tokens of authority. Capability-mode processes have access only to explicitly delegated rights as the use of OS global namespaces is denied.



However, compartmentalization scalability – utilization of increasing numbers of sandboxes – is constrained by performance and programmability limitations of current hardware and software. Today's CPU instruction set architectures (ISAs) reflect a 1990s design consensus conflating virtualization and protection, limiting protection scalability. Compartmentalizing applications using IPC-linked processes also introduces distributed systems programming problems for local applications.

Current systems are exposed to greater threats, demanding dramatically increased use of compartmentalization, placing strain on protection and programming models designed for less risky workloads. CTSRD is developing clean-slate technologies to support large-scale deployment of compartmentalization for the first time.

Peter G. Neumann, Robert N. M. Watson, and Simon W. Moore Jonathan Anderson, Ross Anderson, David Chisnall, Nirav Dave, Brooks Davis, Rance DeLong, Khilan Gudka, Steven Hand, Asif Khan, Myron King, Ben Laurie, Patrick Lincoln, Anil Madhavapeddy, Ilias Marinos, Andrew W. Moore, Alan Mujumdar, Prashanth Mundkur, Steven J. Murdoch, Robert Norton, Philip Paeps, Michael Roe, John Rushby, Hassen Saidi, Muhammad Shahbaz, Stacey Son, Richard Uhler, Philip Withnall, Jonathan Woodruff, Bjoern A. Zeeb

Capability Hardware Enhanced RISC Instructions (CHERI)

CHERI provides a fine-grained protection model within address spaces, complementing existing virtual-memory based process models by providing efficient and more programmable support for application compartmentalization:

We have developed two 64-bit CHERI prototypes, synthesizable for Terasic DE-4 and NetFPGA 10G platforms. Our CheriCloud facility allows remote access to CHERI systems able to run both conventional and adapted software. CheriCloud will be available to early adopters in the DARPA CRASH and MRC communities.

Reference compiler/toolchair

Reference operating system

Reference hypervisor

Hardware research stack Hardware simulation/ implementation substrates

Capability-enabled Clang and LLVM

We have extended the Clang/LLVM compiler suite to generate CHERI ISA instructions based on new C-language annotations and Internal Representation (IR) intrinsics. We add a __capability qualifier for pointers, causing CHERI instructions to be generated, inferring bounds checks and permissions from type information such as arguments to allocation functions and const. These are then dynamically enforced. MIPS and CHERI-aware functions may be interleaved seamlessly allowing a gradual migration. Memory capabilities form the type-safety foundation for shared memory between protection domains. # Call malloc() Bounds checking

__capability int *buffer =

fillArray(buffer, size);

nt retVal = buffer[size];

Capability tagging an // If this isn't a v if (!__builtin_cheri

// const is enforce __capability const [.]

// This will abort fillArray((__capabi __builtin_cheri_get

We provide a set of builtin functions for manipulating capabilities directly in C and enforce properties at run time. This example inspects the length and passes that as an argument, and has a quick-exit path if handed an invalid capability.

http://www.cl.cam.ac.uk/research/security/ctsrd/cheri.html

• Uses a reduced instruction set computer (RISC) approach, providing tools for compiler and operating system writers while minimizing hardware complexity. • Targets low-level software TCBs: OS kernels, language runtimes and web browsers, as well as high-risk data processing such as video decoding. • Allows simultaneous implementation of different security models, reflecting diverse OSs, programming languages, and application requirements. • Implements a hybrid capability model supporting current software side-by-side with components employing fine-grained compartmentalization.







The CHERI-enabled C compiler will automatically set size information for capabilities constructed from dynamic and static allocations and allow automatic dynamic range checking.

5 ••••••••••••••••••••••••••••••••••••	# Get the tag bit
	CGetTag \$1 , \$ c1
and permissions	andi \$1, \$1, 1 # If it's zero, skip to the return
<pre>valid capability, do nothing ri_get_cap_tag(array)) return;</pre>	<pre>beq \$1, \$zero, \$BB1_2 nop # Clear the write permission</pre>
ed with capability protections int *x = array;	<pre>ori \$2, \$zero, 65495 candperm \$c1, \$c1, \$2 # Load the address of the fillArray() function</pre>
<pre>at run time: ility int*)x, t_cap_length(x));</pre>	<pre>ld \$25, %call16(fillArray)(\$1) # Get the capability length for the second argument cgetlen \$2, \$c1 # Call fillArray icln = \$25</pre>
	jalr \$25

Security Oriented Analysis of Application Programs (SOAAP)

Experience with Capsicum shows that adapting programs for compartmentalization is difficult, leading to problems with correctness, performance, complexity, and critically, security. SOAAP is a set of semi-automated techniques to assist programmers with compartmentalization.

Compartmentalization hypotheses are explored through source-code annotations describing sandboxing strategy (e.g., sandbox creation, rights delegation, and RPC forwarding). Security goals and properties (such as information flow constraints and past vulnerabilities) as well as acceptable performance overhead can also be labeled in program source code.

SOAAP uses static and dynamic analysis to engage developers in interactive dialogue, identifying potential correctness bugs (e.g., URL-specific sandbox data inconsistencies), security breaches (e.g., information leaks), expected performance and the impact of software supplychain trojans. SOAAP builds on Clang and LLVM.

	1 2 3 4	<pre>soaap_var_read("compress") int some int some_other_flag = 1; void main() {</pre>
	5 23	<pre>sodap_create_sanabox("compress") some_flag = get_option_from_cmd_line</pre>
	24 25 26	<pre>gz_compress(in, out); }</pre>
	27 28 29 30 31	<pre>soaap_sandbox_persistent("compress") void gz_compress(int ifd, int ofd) { if (some_flag && some_other_flag) { </pre>
	32	} }
Past vulnerabilities		
	1 2 3	<pre>soaap_sandbox_ephemeral("parser") void parse(soaap_read_fd int ifd, D</pre>
	4 5	if () { soaap_vuln_pt("CVE-2005-ABC");
	10	···· }
	13 14 15	} soaap_vuln_fn("CVE-2005-DEF")
	16 17 18	<pre>void not_sandboxed() { }</pre>
	Cor	nfidentiality
	1 2	<pre>soaap_classify("secret") char* serv</pre>
	3	<pre>void main() { </pre>
	20 21 22 23	<pre>while () { accept_connection(); } }</pre>
	24 25 26	<pre>soaap_sandbox_persistent("session") void accept_connection() {</pre>
	27 28 29	<pre>soaap_private char session_key[10 compute_session_key(session_key, se }</pre>
		formance
	Peri	
	Peri	<pre>soaap_sandbox_ephemeral("decoder")</pre>
	Peri	<pre>soaap_sandbox_ephemeral("decoder")soaap_sandbox_overhead(12) int jpg_decode(char* buf.</pre>
	1 2 3	<pre>soaap_sandbox_ephemeral("decoder")soaap_sandbox_overhead(12) int jpg_decode(char* buf,soaap_data_in int le</pre>

http://www.cl.cam.ac.uk/research/security/ctsrd/soaap.html

Site-specific sandbox

Site-specific sandbox



peated SOAAP iteration as program and hypotheses are refined



Temporally Enhanced Security Logic Assertions (TESLA)

TESLA allows programmers to describe temporal properties of security-critical software and validate them at runtime.

Programmers describe these properties with inline assertions or explicit finite-state automata. Both forms of TESLA specification are written in C, referencing names from surrounding scopes and exploiting the compiler's type checker. Both are converted to a TESLA intermediate representation (IR), allowing other languages to target the TESLA backend as well.



In the example below, the programmer asserts that the X.509 certificate passed to the use_cert function has been properly verified. In fact, however, the certificate verification code suffers from the vulnerability in CVE-2008-5077: an OpenSSL error code has been misinterpreted as success.





In this example, TESLA observes the `main()` entry event, so it creates an automaton instance and moves it from state 0 to state 1. However, it does not observe the $X509_STORE_CTX_init() == 1$ event (this is the cause of the verification flaw), so when the `NOW` event occurs, TESLA cannot find an automaton instance named $(cert=0x7fda614147c0, \star, \star, \star)$. The certificate has not been verified!

http://www.cl.cam.ac.uk/research/security/ctsrd/tesla.html











Mr Rance













Mr Brook







Mr Jonathar Woodruff



Members of the CTSRD team and its external oversight group at our May 2011 review meeting in Cambridge, UK Joe Stoy (Bluespec), Jonathan Woodruff (Cambridge), Ben Laurie (Google), Ross Anderson (Cambridge) Virgil Gligor (CMU), Philip Paeps (Cambridge), Li Gong (Mozilla), Peter Neumann (SRI) Simon Cooper, Michael Roe (Cambridge), Robert Watson (Cambridge), Howie Shrobe (DARPA) Steven Murdoch (Cambridge), Sam Weber (NSF), Jonathan Anderson (Cambridge), Simon Moore (Cambridge) Anil Madhavapeddy (Cambridge), Dan Adams (DARPA), Rance DeLong (LynuxWorks). Jeremy Epstein (SRI), Hassen Saidi (SRI)



Many Cambridge members of the CTSRD team meeting for an August 2012 project photo in Cambridge, UK Jonathan Woodruff, Richard Clayton, Jonathon Anderson, Michael Roe, Ross Anderson, David Chisnall, Robert Watson Khilan Gudka, Robert Norton, Simon Moore





Approved for public release. This research is sponsored by the Defense Advanced Research Projects Agency (DARPA) and the Air Force Research Laboratory (AFRL), under contract FA8750-10-C-0237. The views, opinions, and/or findings contained in this article/ presentation are those of the author/presenter and should not be interpreted as representing the official views or policies, either expressed or implied, of the Defense Advanced Research Projects Agency or the Department of Defense.







