



CHERI tablet and rack-mount CheriCloud 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, programmer-friendly, 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, running current C-language operating systems and applications while supporting gradual adoption of novel protection features in both critical Trusted Computing Bases (TCBs) and high-risk software components.



### Capsicum and the application compartmentalization motivation

Programmers are turning to application compartmentalization to mitigate inevitable vulnerabilities: software is decomposed into sandboxes in accordance with the principle of least privilege. As granularity increases, individual sandbox rights decrease. Only the rights of compromised components are leaked, forcing attackers to exploit more vulnerabilities to accomplish the same goals.

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.

### Capability Hardware Enhanced RISC Instructions (CHERI)

CHERI provides fine-grained protection within address spaces, complementing virtual-memory based processes to efficiently support compartmentalization:

- 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.

We have developed two CHERI prototypes, synthesizable to FPGA. CheriCloud allows remote access to systems running conventional and adapted software. BERI Open Systems CIC, a non-profit UK company, will open source CHERI in 2014.

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, Alex Horsman, Jong Hun Han, Asif Khan, Myron King, Ben Laurie, Patrick Lincoln, Anil Madhavapeddy, Ilias Marinos, Dr Theo A. Markettos, Ed Maste, Andrew W. Moore, Alan Mujumdar, Prashanth Mundkur, Steven J. Murdoch, Robert Norton, Philip Paeps, Michael Roe, Colin Rothwell, John Rushby, Hassen Saidi, Muhammad Shahbaz, Stacey Son, Richard Uhler, Philip Withnall, Jonathan Woodruff, Bjoern A. Zeeb

## CheriBSD hybrid capability OS

CheriBSD is an adaptation of the FreeBSD operating system to utilize features of the CHERI ISA. CheriBSD relies heavily on CHERI's hybridization support, allowing us to deploy in-process memory protection and sandbox incrementally. FreeBSD has been extended to:

% ssh ctsrd@cheritest.sec.cl.cam.ac.uk Last login: Thu Nov 14 01:54:08 2013 from c0188.aw.cl.cam.ac.uk FreeBSD 11.0-CURRENT (CHERI_DE4_SDROOT) #4 011d9a4(master)-dirty: Mon Jan 6 13:19:44 GMT 2014 rt@c0188:~% slogin rnw24@cheritest.sec.cl.cam.ac.uk										
% procstat -RX 1	1180									
PID COMM	CLASS	METHOD	INVOKE	FAULT	LMIN	LMAX	SMIN	SMAX	SMEAN	SMEDIAN
11180 cheritest	cheritest-helper.bin	md5	5	0	11133	56396	11133	56396	22383	11831
11180 cheritest	cheritest-helper.bin	abort	2	2	3048	3288	3048	3288	3168	3168
11180 cheritest	cheritest-helper.bin	puts	3	0	419754	598131	419754	598131	508887	508777
11180 cheritest	cheritest-helper.bin	syscall	3	0	6836	7581	6836	7581	7228	7269
11180 cheritest	cheritest-helper.bin	divzero	4	4	2915	3188	2915	3188	3050	3049

In the coming months, we plan to implement further CheriBSD features including deploying CHERI memory protection through more of the CheriBSD userspace, and deploying additional sandboxing around the system.

# Capability-enabled Clang and LLVM

We have extended the Clang/LLVM compiler to generate CHERI ISA instructions based on new C-language annotations and Internal Representation (IR) intrinsics. We add a \_\_\_capability pointer qualifier, triggering CHERI instructions generation, inferring bounds checks and permissions from type information (e.g., arguments to allocation functions and const). These are dynamically enforced. MIPS and CHERI functions interleave seamlessly allowing gradual migration. Memory capabilities are the type-safety foundation for shared memory between protection domains.

Bounds checking
<pre>#include <capabili< pre=""></capabili<></pre>
<pre>// The compiler au capability int *bu (capability int</pre>
<pre>// Size can be com // made explicit so // non-capability-o fillArray(buffer,</pre>
<pre>// This overflows int retVal = buffe</pre>
Capability tagging a
<pre>// If this isn't a if (!builtin_che</pre>
<pre>// const is enforce capability const is</pre>
// This will abort

• Maintain CHERI registers for user threads • Selectively utilize CHERI memory

protection via CHERI Clang/LLVM Implement CCall/CReturn fast exceptions:

object-capability invocation

• Implement a CHERI "trusted stack"

tracking object-capability invocation

• Recover from sandbox faults, such as



Hybrid code blending general-purpose registers and capabilities

High-assurance "pure" capability code

Per-address space memory management and capability executive

memory errors, returning control to the invoking application (or sandbox) • Provide a libcheri(3) API allowing applications to create and invoke sandboxes • Implement monitoring extensions to the procstat(1) tool to track sandbox state • Support CHERI debugging extensions for CHERI LLDB

• Support on-board peripherals such as Ethernet, display, flash, and SD card



### Application compartmentalization on CHERI

With increasingly mature CHERI processor prototypes, CHERI Clang/LLVM, and CTSRD CheriBSD, we are turning our attention to SRI International and the University of Cambridge deploying CHERI features around the OS and Collaboration spans historically siloed research are: Security, CPU architecture, operating systems, application. In November 2012, we compilers, programming languages, formal methods Clean slate design violates conventions in exchange for demonstrated a bespoke presentation dramatic security improvements package, CheriPoint, able to perform apability-based compartmentalization mitigate known and unknown classes of vulnerabilities granular sandboxing when rendering composed slide decks from many sources. We are now applying CHERI features in off-the-shelf applications such as tcpdump, experimenting with memory protection and fine-grained sandboxing.

tcpdump -i atse0 | head -20 tcpdump: verbose output suppressed, use -v or -vv for full protocol decode istening on atse0, link-type EN10MB (Ethernet), capture size 65535 bytes , val 91597245 ecr 1123802448], length 0 win 1040, options [nop.nop.TS val 1123802478 ecr 91597245], length 116 options [nop.nop.TS val 91597275 ecr 1123802478], length 0

tcpdump is a widely used packet capture and interpretation tool; it is able to render and analyse packet sequences including low-level network characteristics (e.g., Ethernet/IP/TCP headers) and application-level features (e.g., NFS RPC decoding). With a large library of *protocol dissectors* able to render packets from a broad range of protocols, tcpdump is subject to a large attack surface. Further, tcpdump is also frequently used in security-sensitive environments: analysing traffic off the Internet, and when diagnosing suspicious traffic or in-progress attacks.

cheri\_tcpdump is an adaptation of tcpdump using the CHERI model: memory protection, fine-grained compartmentalization, and incremental deployability. We are using tcpdump to explore dynamic compartmentalization: adaptive adjustment of performance investment in security responding to changing security parameters.

## 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 techniques to assist 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 (e.g., information flow constraints and past vulnerabilities) as well as acceptable performance overhead are labeled in source code.

SOAAP engages the developer in a dialogue, identifying potential bugs (e.g., data inconsistencies), security breaches (e.g., information leaks), expected performance, and the impact of software supply-chain trojans. Past vulnerabilitie

<pre>1soaap_sandbox_ephemeral("parser")</pre>
<pre>2 void parse(soaap_read_fd int ifd, D</pre>
3
4 if () {
<pre>5soaap_vuln_pt("CVE-2005-ABC");</pre>
• • •
10 }
•••
13 }
14
15soaap_vuln_fn("CVE-2005-DEF")
16 void not_sandboxed() {
17
18 }









# **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.

void
use_cert(X509 *cert)
{
<pre>#ifdef TESLA</pre>
TESLA_WITHIN(main, previously(
<pre>X509_STORE_CTX_init(ANY(ptr), ANY(ptr), cert, ANY(ptr)) =</pre>
X509_verify_cert(ANY(ptr)) == 1
));
#endif
<pre>/* use the certificate */</pre>
}



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, \*, \*, \*). The certificate has not been verified!



























Dr Theo A Markettos









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 (Cambridg Virgil Gligor (CMU), Philip Paeps (Cambridge), Li Gong (Mozilla), Peter Neumann (SRI) non Cooper, Michael Roe (Cambridge), Robert Watson (Cambridge), Howie Shrobe (DARPA en 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)



Members of the CTSRD and MRC2 teams meet for our August 2013 annual meetings in Cambridge, UK.

Ed Maste, Prashanth Mundkur, Steven Murdoch, Jong Hun Han, Hassen Saidi, Khilan Gudka, Colin Rothwe Peter G. Neumann, Malte Schwarzkopf, Brooks Davis, Nirav Dave, Jonathan Woodruff, Bjoern Zeeb, Robert Norton, David Chisnall, Alan Mujumdar, Alex Horsman

Andrew Moore, Simon Moore, Robert Watson



**UNIVERSITY OF CAMBRIDGE** 

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.