

CHERI Alliance: CHERITech'26

Keynote - With CHERI, can we afford to ignore hardware security?

J. Bormann
Siemens EDA, Munich, Germany



Abstract

In response to the evolving geopolitical landscape, governments are increasingly driving cybersecurity through regulatory frameworks, such as the Cyber Security Act, and by incentivizing research through targeted funding. Siemens EDA occupies a pivotal role at the intersection of academic research and semiconductor manufacturing. We fulfill this role by actively contributing to the cybersecurity community—partnering with the CHERI Alliance and shaping standards like the Common Weaknesses Enumeration (CWE) and about security collaterals for third-party IP at IEEE. Furthermore, our verification portfolio leverages formal methods to identify deeply embedded malicious modifications in hardware, addressing a critical risk vector in CHERI-based system design.

Biography

Jörg Bormann is Product Manager for Advanced Verification at Siemens EDA. He is responsible for Siemens EDA's formal Trust and Security verification technology and the GapFree solution, that formally verifies all circuit functionality. Previous employments include Infineon, Intel and OneSpin Solutions.

Jörg Bormann received in 2009 his Dr.-Ing. from the University of Kaiserslautern after an extra-occupational graduation on the GapFree verification approach. His professional interest includes new approaches for the verification of the interaction of firmware and hardware of embedded systems.

References

Progress of the Alliance

M. Eftimakis
CHERI Alliance, Cambridge, United Kingdom



Abstract

This talk will discuss the progress of the CHERI Alliance since its official launch in November 2024.

Biography

Mike Eftimakis has an extensive background in the semiconductor and electronics industry with 30 years in senior technical and business roles. He has a rich history of innovation with companies like VLSI Technology, NewLogic, Arm and Cudasip, and he started-up and led his own company. where he played pivotal roles in advancing technology and business strategies. His expertise ranges from chip design engineering and system architecture to product and company management, marketing and strategy, making him a key contributor to the growth and success of microelectronics organisations.

In parallel to his VP Strategy and Ecosystem role at Cudasip, he is the Founding Director of the CHERI Alliance, an industry association dedicated to promoting CHERI technology. This technology addresses the root causes of most current cyberattacks, contributing to a safer and more trustworthy World.

References

Opportunities for CHERI in the Energy Industry

J. Cooper
Hildebrand Technology Limited, London, United
Kingdom



Abstract

This talk sets out the commercial and regulatory case for CHERI-enabled devices in the GB energy market. It demonstrates that CHERI is not merely a security improvement, it is driven by converging regulatory pressure from the EU Cyber Resilience Act (CRA), the UK Product Security and Telecommunications Infrastructure (PSTI) Act, and the Network and Information Systems (NIS2) in conjunction with the operational economics of running connected device fleets over decade-long deployment horizons.

The cost of not adopting CHERI expressed through breach liability, compliance overhead, firmware update logistics, and constrained product architectures is increasingly more significant than the cost adopting it. After 15 years of development CHERI production silicon is available today, enabling energy sector device manufacturers to build hardware-enforced memory safety into their products. The next step is to demonstrate that certified, commercially viable products can be built on this foundation, showing that CHERI-secured devices can achieve cybersecurity certifications such as CPA more quickly and at lower cost than conventional alternatives.

As part of the CHERI Adoption programme, Hildebrand is bringing to market two fully certified metering products: the cSAPC, a certified smart metering device and the cCOP11 a compliant asset meter enabling participation in the GB flexibility market. All will be built on SCI Semiconductor's ICENI silicon. This will be the first range of products commercially available that are UK-made, CHERI-secured and ready for deployment.

Biography

Joshua Cooper is CEO and co-founder of Hildebrand, a UK technology company specialising in smart energy data and connected technologies. He has extensive experience using real-world energy data to help consumers, businesses and policymakers better understand energy usage.

He has been involved in a range of energy innovation projects exploring how smart meter data can support energy efficiency, heat decarbonisation and the transition to low-carbon technologies. His work includes developing approaches that use actual household energy consumption to understand heating demand and model the impact of energy efficiency measures.

References

Core to edge IoT security with CHERI

P. Cox
UM Labs, London, United Kingdom



Abstract

UM Labs, a member of the CHERI Alliance, is a UK based cybersecurity R&D company specialising in securing real-time communication on IP networks. Our software spans cloud-based core services and edge devices running as embedded systems.

This session reviews UM Labs' experience in porting Unicus® IoT, our security platform to CHERI. Unicus protects IoT services running Zigbee, MQTT and other protocols using a Core/Edge architecture. Unicus Edge connects local IoT sensors and actuators and provides an authenticated and encrypted uplink to Unicus Core which provides provisioning and management functions.

UM Labs have ported both the Core and Edge components of our IoT security platform to CHERI. This process used both hardware and virtualised targets enabling verification of the code at both compile time and run time. Problems identified in either phase of development and testing can be addressed, improving the quality and security of the delivered software. This process reduces the risk that latent flaws in memory management and code separation cause security violations in production. Using CHERI to build a complete commercially deployable end-to-end network security service demonstrates the maturity of the technology and shows the value of CHERI in any development project.

Biography

Peter Cox is the CEO and founder of UM Labs, a UK based R&D Company specialising in developing cyber security technology for real-time communications. Peter has over 30 years of cyber security experience and was the co-founder of one of the first commercial firewall companies.

At UM Labs Peter designed the Unicus® security platform for real-time communication security covering voice and video telephony, streaming video, cloud extended reality and IoT. Unicus is widely installed in defence, government and commercial networks where it secures communications for a European National Defence Force.

References

Building a Network Security Product on CHERI: Early Insights

M. Atkins
Mission Critical Applications Limited, Bath, United
Kingdom



D. Atkins
Mission Critical Applications Limited, Bath, United
Kingdom



Abstract

Routers occupy a uniquely privileged position in every network — they see, and can redirect, almost every packet that passes through. Their management interfaces are consistently among the most attacked surfaces in networking, and 2026 has already brought a string of serious disclosures affecting major vendors. Compromise of a router's management interface can have consequences well beyond the device itself, across enterprise, telecoms and national-infrastructure networks alike.

We are a UK SME building a hardware-enforced security appliance to address this problem, using CHERI on Arm Morello with a defined path to CHERI-RISC-V. Our premise is simple: a device trusted to guard a network's most sensitive interface must itself be measurably more robust than the system it protects.

In this talk we will share why we believe CHERI's memory-safety and compartmentalisation guarantees are particularly well suited to problems of this kind, the practical realities of building a commercial product on a CHERI-compatible platform, and the open question of what should count as credible evidence that a security product is trustworthy, rather than merely marketed as such. We're also keen to compare notes with others in the CHERI community bringing CHERI-based products from research prototype toward commercial deployment, particularly around platform choices and evidencing trust to non-technical stakeholders.

This is an early-stage project, and we see CHERITech'26 as a chance to engage the community as our work develops.

Biography

Dr Martin Atkins

Dr Martin Atkins is Technical Director at Mission Critical Applications Limited, a UK SME specialising in high-integrity, safety-critical systems across aerospace, telecoms, automotive and defence. He holds a DPhil in Computer Science and a BA in Mathematics from the University of York. His interest in hardware capability systems predates CHERI itself, dating back to post-doctoral work on the Ten15 and Flex capability architectures at the Royal Signals and Radar Establishment. He was technical lead on MCA's work within the Digital Security by Design (DSbD) Technology Access Programme and the subsequent DASA-funded "CHERI within Defence and Security" programme, exploring seL4 on Arm Morello. He specialises in rapid prototyping, test harnesses and open-source safety-critical development.

Dr Divya Atkins

Dr Divya Atkins is Managing Director of Mission Critical Applications Limited, a UK SME specialising in high-integrity, safety-critical systems across aerospace, telecoms, automotive and defence. She holds a BSc from Mount Carmel College, Bangalore University, a BE in Computer Science from the Indian Institute of Science, an MSc in Computing from Balliol College, Oxford, and a DPhil from the University of York. She has led MCA for over 25 years, with a research background in formal methods and industrial software engineering standards. She was project lead on MCA's work within the Digital Security by Design (DSbD) Technology Access Programme and the subsequent DASA-funded "CHERI within Defence and Security" programme, with a particular interest in computer and network security.

References

Beyond the tech and what it takes to create a market for CHERI

M. Ravishankar
PA Consulting, London, United Kingdom



Abstract

CHERI has reached a pivotal point in its technology journey. After over a decade of research and technical development, the focus is increasingly shifting from proving the technology works to understanding what conditions are required for adoption at scale. As organisations invest in CHERI-enabled products, toolchains and hardware platforms, the challenge is moving from technology development towards ecosystem growth and market creation.

The presentation will share lessons from the UK's CHERI Managed Deployment Programme which was established to not only demonstrate CHERI in real-world use case but better understand what it takes to accelerate adoption and so pave the way for future waves. Bringing together technology providers, product developers, system integrators, end users and government, the programme provides an important view of the opportunities and challenges involved in translating a promising technology like CHERI into a sustainable market.

Drawing on experiences from the ongoing deployments, the presentation will explore what we've learned so far about building the conditions for adoption. This includes understanding how different sectors value CHERI, what it takes to engage OEMs and product vendors, how assurance can support adoption, and the role that hardware, software and supply chain dependencies play in deployment decisions.

One of the key insights emerging from the programme is that there is a significant gap between a technology being technically viable and a market being ready to adopt it! Organisations developing CHERI-enabled products must also build confidence amongst customers, partners and supply chains that the technology delivers meaningful value. Across the programme, we have observed that building a CHERI-enabled product does not automatically create demand for it. Understanding how security benefits are recognised, valued and incorporated into purchasing and deployment decisions is emerging as an important area of learning.

The programme is also providing early insights into some of the factors likely to shape the next phase of adoption. These include the role of assurance in building trust and reducing adoption risk, the importance of mature hardware and software supply chains, and the challenge of generating sufficient market pull to encourage OEMs and product vendors to incorporate CHERI into future product roadmaps. Experience to date suggests that successful adoption depends not only on deployable technology, but on the wider ecosystem of evidence, partnerships, products and market signals that sit around it.

The talk will reflect on how these insights are shaping the programme and discuss what they mean for the future of CHERI. It will take a cross-programme perspective on the ecosystem, partnerships and market

conditions needed to move CHERI from promising technology towards deployable products, broader adoption and sustainable commercial opportunities.

Biography

Manu Ravishankar is a Principal Consultant at PA Consulting, where he leads innovation, strategy and delivery programmes across cybersecurity, digital infrastructure and critical national infrastructure. He led the design and establishment of the UK's CHERI Secure by Design Managed Deployment Programme and now serves as its Programme Delivery Lead, bringing together government, industry and technology partners to support the deployment and adoption of CHERI-enabled technologies. His work focuses on scaling emerging technologies from research and development into real-world deployment and commercial adoption.

References

High performance CHERI tag controller

A. Joannou
University of Cambridge, Cambridge, United
Kingdom



Abstract

CHERI's memory safety relies on its capabilities being unforgeable. A CHERI system provides this unforgeability guarantee by using a tagged memory: for each memory location that can hold a capability (128-bits), there exists an out-of-band tag to indicate raw data or valid CHERI capability. When a tag is set, it is safe to use the corresponding 128-bit data as a capability. Every memory write must carry a tag (1 if a valid capability is written, 0 for any other write) to ensure no capability can be crafted without obeying strict capability derivation rules. Current memory systems do not typically support tagged memory, and therefore need extending to support CHERI.

The AXI tag controller aims to enable integration of CHERI capability tags with a standard memory system. It receives a tagged request which uses the AXI user field to carry the tag associated with a capability carried in the data field, and splits it into two requests, one for tag, and one for data. Neither of these new requests requires tagged memory, and they can therefore be serviced by conventional memory. The data request can be forwarded unchanged (except for the user field no longer being present), and the tag request can be aimed at a dedicated region of the standard memory, only accessed by the tag controller as a tag store.

The tag controller provides some configuration flexibility to map the covered region and the tag store region, prevents access to the later, and uses it to back a hierarchical tag table consisting of root and leaf tag bits. Together with caching, using a hierarchical approach minimises required bandwidth for tag traffic, and ensures very low overheads for accesses to large regions of memory without tags. It uses a separate cache for the root table and the leaf table, allowing for low lookup latency on hit.

The system verilog IP implementing the AXI tag controller leverages the HPDCCache as its tag caches, and offers parameterizable tag table grouping factor and cache sizes as well as capability sizes should alternative CHERI format be explored. It can be used to enable CHERI systems to integrate with otherwise tag-oblivious memories, and is in use in our CVA6-CHERI work.

Biography

Alexandre Joannou is a senior researcher at the University of Cambridge. He joined the CHERI research team in 2013 for his PhD and has been contributing to the project since then at a architectural and micro-architectural level.

References

CheriBSD on RISC-V VP++: A fast virtual prototype for latest RVY 0.9.9

A. Hinterdorfer

Johannes Kepler University (JKU) Linz, Institute for
Complex Systems (ICS), Linz, Austria



Abstract

We present RISC-V VP++, an open-source virtual prototype for exploring and evaluating RISC-V systems with CHERI support. While the platform has supported earlier CHERI ISA versions, it has now been brought up to the latest RVY specification (0.9.9), enabling experimentation with the current state of CHERI RISC-V. The updated implementation covers all architectural changes introduced by RVY0.9.9 and has been validated against QEMU using TestRIG. The differential comparison between the two independently developed platforms (QEMU vs. RISC-V VP++) provides strong confidence in their agreement on the implemented instruction semantics, capability handling, privilege mechanisms, memory behavior and CHERI protection. These validations not only strengthen RISC-V VP++, but also QEMU's confidence in conformance with the latest RVY specification.

In contrast to high-level emulators such as QEMU, RISC-V VP++ is designed as an extensible transaction-level virtual prototype. Its SystemC/TLM-based architecture exposes ISS and platform behavior at a level suitable for architectural experimentation, instrumentation, hardware/software co-simulation and rapid evaluation of new extensions.

At the same time, its simulation performance is sufficient to boot the full-scale, general-purpose CheriBSD operating system and interact with it at near real-time speeds, making full-system experimentation practical rather than limited to small bare-metal workloads. Running CheriBSD demonstrates that the implementation extends well beyond individual instruction semantics. This confirms the entire hardware/software stack agrees on capabilities usage, virtual memory and trap handling, when firmware is exercised in a realistic full-system environment.

We demonstrate RISC-V VP++ running the latest CheriBSD on RVY 0.9.9, booted through OpenSBI. The demonstration includes interactive use of the operating system and examples of CHERI capability enforcement. It also highlights how the virtual prototype can be used to inspect and modify architectural behavior in ways that are difficult or impossible to achieve on real hardware. This includes selectively relaxing CHERI protection at runtime for controlled debugging. This feature allows revisiting the work on counterfactual execution to evaluate the effectiveness of CHERI protection, now using RVY 0.9.9 in a full operating-system environment [1].

By combining support for the latest RVY specification, validated architectural behavior, and full-system CheriBSD execution, RISC-V VP++ provides a practical platform for cutting-edge experiments with the rapidly evolving CHERI RISC-V ecosystem. Because the complete platform runs in software on conventional host systems, researchers can experiment with latest CHERI RISC-V without access to dedicated CHERI hardware, FPGA infrastructure, or an RTL simulation environment. As an open-source project, RISC-V VP++ therefore significantly lowers the barrier to entry for researchers and potential CHERI adopters.

References:

[1] Distinguishing Exploit Failure from Effective CHERI Protection on RISC-V, Andreas Hinterdorfer, Manfred Schlägl, Daniel Große, RISC-V Summit Europe 2026

Biography

Andreas Hinterdorfer is a PhD student and research assistant at the Institute for Complex Systems (ICS), JKU Linz, under the supervision of Prof. Daniel Große. He received his MSc in Electronics and Information Technology with distinction in 2025 and subsequently worked as embedded software engineer at blue-zone GmbH Austria. His research focuses around hardware/software co-simulation, virtual prototyping and processor verification. He began working on CHERI during his master thesis and continues to research CHERI-enabled RISC-V systems, with a particular focus on virtual prototypes and validation.

References

Mapping the ETSI Memory-Safety Technical Standard to CHERI

G. Jenkinson
Capabilities Limited, Cambridge, United Kingdom



CAPABILITIES
LIMITED

Abstract

Memory unsafety remains the leading cause of exploitable software vulnerabilities across systems software. While mitigations range from probabilistic software and hardware mechanisms to memory-safe languages and formal verification, the security community has historically lacked a common, vendor-neutral taxonomy to evaluate and compare these approaches.

To address this gap, ETSI TC CYBER is developing TS 104 198, establishing formal definitions for memory-access models, memory-safety properties, enforcement semantics, completeness, and assurance levels. This talk provides a structured, technical exploration of the standard and its mapping to CHERI architectural capabilities and real-world CHERI platforms:

- Introduction & Motivation: An overview of ETSI TS 104 198, establishing why a technology-neutral vocabulary for memory safety is essential for industry procurement and technical roadmaps.
- Introduction to ETSI TS 104 198: A deep dive into the abstract memory-access models and core memory-safety properties defined in the standard, including reference, spatial, and temporal safety.
- Mapping CHERI Hardware Capabilities to ETSI Properties: A detailed technical mapping showing how CHERI's unforgeable capabilities, approximate bounds, and revocation mechanisms satisfy normative criteria for complete, deterministic memory safety. Known limitations of CHERI against the standards properties will also be presented.
- Broader Industry Impact & Future Work: Overview on the envisaged usage of the standard to develop and publish transparent product roadmaps. Brief updates on current work on the draft specification and opportunities to participate in ETSI standardisation activities.

Biography

Graeme Jenkinson serves as Director of Applied Technology at Capabilities Limited, where he has led various projects, such as web-service compartmentalisation and porting the V8 JavaScript Engine and Chromium web-browser to CHERI C/C++ on ARM Morello. Bringing over two decades of expertise across industry and academia, Dr. Jenkinson has previously managed a security research domain at BAE Systems, engineered confidential computing platforms at Apple, and contributed to the DARPA Transparent Computing program while at the University of Cambridge. As an active member of the ETSI TC CYBER

memory safety working group (DTS/CYBER-00165 / TS 104 198), he has made significant contributions to specifying the memory-access models, formal property definitions, and vendor-neutral terminology.

References

CHERI Specification and Standardisation Update for RV64Y, RV32Y and CHERIoT

T. Kurd
Codasip, Cambridge, United Kingdom



Abstract

This talk will give the status of the standardisation process in RISC-V.

It will cover the RV64LYA base ISA for Linux-capable application cores, and also RV32LYA for general purpose micro-controllers and embedded cores, and the highly specialised CHERIoT format.

It will also include details of future extensions which are planned to improve the performance and/or security of RV64LYA cores.

The talk will be written shortly before the deadline to relect the current status.

The intention is to announce a ratification date of the RV64LYA base ISA and the relevant profile in the spring of 2027, and that the specification is already in a frozen state with upstreamed toolchains, simulators and tests etc. available for public use.

Biography

Tariq is Chief Architect and is a Distinguished Engineer at Codasip. He has over 30 years of experience in the silicon industry, and has worked for major companies such as NVIDIA, Broadcom, Huawei, and STMicroelectronics. He is an active participant in the RISC-V standards body, and a TSC member. He has ratified multiple RISC-V extensions. He has spent 5 years working on CHERI and aims to bring memory safety to mainstream computing.

References

A Standard for CHERI Capability Transport over AXI and CHI

B. Fletcher
CHERI Alliance, Cambridge, United Kingdom



Abstract

CHERI extends conventional memory safety by adding architectural capabilities: values that carry authority to access memory. For the purposes of system interconnect, a capability can be thought of as a combination of data and a hardware-maintained tag bit. The tag bit is small, but it is vital for maintaining the memory safety guarantees that CHERI offers. It indicates whether the associated data should be treated as a valid capability or as ordinary data. If tag state is lost, corrupted, or incorrectly created while data moves through a system, the security properties of CHERI can be weakened or broken.

This talk introduces the “CHERI Tag Bit Transport Overlay Standard”, a developing specification for transporting capabilities across system fabrics and bus protocols. The goal of the standard is to provide a framework for moving capability data and tag state over common protocols such as AXI and CHI. By defining how capability data and its associated tag state are transported alongside ordinary read and write data, the standard aims to make successful implementation of CHERI technology easier and improve interoperability between independently developed CHERI-aware IP blocks.

The talk will follow the structure of the standard. It begins with the protocol-independent issues that need to be considered when transporting capabilities around a system, and the rules that must be followed to maintain the guarantees offered by CHERI. This includes preserving data/tag atomicity, defining what must happen if that atomicity is violated, maintaining tag validity as capabilities move between components, and considering what happens when the width of a system bus is different from the width of a capability.

It then moves on to the protocol-specific mappings. For AXI, the standard defines how tag bits are transported using AXI user signaling. For CHI, the standard maps tag transport onto CHI RSVDC signaling and highlights the requirement that these bits must be preserved by the CHI implementation for the overlay to interoperate correctly. TileLink is discussed as future work: it is expected to be added to the standard, but is not currently part of the specification.

Finally, the talk will cover some of the practical system issues that arise when capabilities move through a real SoC, including width conversion, caching, and coherency. These topics matter because intermediate system components may not create capability authority, but they can still weaken the guarantees offered by CHERI if they lose, corrupt, or mis-associate tag state.

The aim of the talk is to explain why the standard exists, what problem it solves, and where review is most valuable. Feedback on the version 0.9 draft is actively sought, with the intention of incorporating review comments into a version 1.0 release in 2027.

Biography

Ben Fletcher is IP Engineering Director at Codasip, where he leads the CPU Core Architecture team and is currently working on architectures for CHERI-enabled out-of-order RISC-V CPU cores.

He is also “CHERI in SoC” Work Group Lead at the CHERI Alliance, where his current work focuses on the practical system-level adoption of CHERI technology, including how CHERI-aware processors, interconnects, memories, and peripherals can be integrated into interoperable SoCs.

Ben has more than 25 years of experience in semiconductor and SoC development, spanning architecture, design, verification, and software. Before joining Codasip, he held engineering leadership roles at Sondrel, an engineering management role at Imagination Technologies, and technical roles at Broadcom and Oxford

Semiconductor.

References

CVA6-CHERI - Ratified Specification Implemented and Vector Support Planned

J. Woodruff
Capabilities Limited, Cambridge, United Kingdom



CAPABILITIES
LIMITED

Abstract

CVA6-CHERI has been under development as an ASIC-ready open-source demonstrator for the ratified RVY CHERI specification for the last two years. The Capabilities Limited team has recently pulled CVA6-CHERI forward to the ratified version of the RVY specification, along with QEMU, LLVM, OpenSBI and CheriBSD, achieving boot on the new specification on FPGA. New features include richer PTE enforcement, including SSTATUS bits for controlling generational revocation in the kernel, and removal of exceptions on capability branches. Timing has improved due to the architecture upgrade, and CVA6-CHERI now normally passes timing at 50MHz on the Genesys2 FPGA. CVA6-CHERI has also completed changes in response to formal verification, and is fielding issues from ASIC consumers as maturity approaches.

We have also begun plans for the next phase of CVA6-CHERI development. The AIVHAI project plans to extend CVA6-CHERI with standard vector and matrix extensions using a hybrid of the PULP Ara interface and the Google Coral vector unit. We will describe our plans in this direction, which include AI-automated formal verification of the components and integrated system.

Biography

Dr Jonathan Woodruff is an Assistant Research Professor with expertise in processor architecture and microarchitecture as well as low-level software optimisation. Specialising in capability processor design, he has pushed into full-system optimisations including cache hierarchy, core timing, and multi-core designs as well as explorations into major security approaches including control flow integrity and private execution.

References

CHERI in OpenTitan: application of memory safety best practice in an open hardware root of trust

J. Orensanz Martinez
lowRISC, Cambridge, United Kingdom



Abstract

OpenTitan is a popular open source hardware root of trust implementation, used in all sorts of applications, including Google's Chromebook devices. The new OpenTitan top-level design Earl Grey 2 will incorporate CHERI. This presentation covers why CHERI is relevant (essential!) to a RoT design, the kind of software attacks and defects it protects from and how they are complementary to the hardware attacks that the current design focuses on. It will also touch on the specific implementation of CHERI in OpenTitan and the estimated impact it will have in terms of area and cost. Finally, it will touch on how to get hold of RISC-V based implementations of CHERI provided by lowRISC: Earl grey 2 and CHERI-Mocha.

Biography

Javier is the CEO of lowRISC, a non-profit with a mission to make open silicon a commercial reality, who runs active projects on secure and memory safe hardware. Previously, he was a VP and General Manager at Arm for over 8 years. Javier has worked in the semiconductor industry for over 25 years.

References

Building Memory Safety into a DMA Engine with CHERI

B. Fletcher
CHERI Alliance, Cambridge, United Kingdom



Abstract

DMA engines read and write memory independently of the CPU, and conventionally rely on an external mechanism - like an IOMMU - to police what they're allowed to touch. History shows that boundary is often incomplete, misconfigured, or simply not yet active when it matters, from early-boot DMA to malicious peripherals like those behind the Thunderclap attacks. CHERI takes a different approach: instead of policing access from outside, it makes the constraint part of what's handed to the device in the first place.

This talk works through what that shift means in practice, building up from first principles. We start with the simplest possible DMA engine - programme source and destination registers, kick it off, wait for an interrupt - and show what changes when capabilities replace raw addresses: how they're delivered to the engine, how the engine constrains its own accesses to what it was given, and how cleanup on completion helps to deliver temporal safety. We then repeat the exercise for a more realistic engine driven by a linked list of descriptors, where the capabilities themselves must be chained safely through memory the engine doesn't fully control.

Finally, we show this isn't just a paper exercise: Codasip is currently developing a CHERI-aware DMA engine internally, supporting both register-based and descriptor-based operation. Ahead of hardware, we've built a QEMU model of the engine, which we're using for early software development and to validate the architecture and software interface before committing to RTL.

Biography

Ben Fletcher is IP Engineering Director at Codasip, where he leads the CPU Core Architecture team and is currently working on architectures for CHERI-enabled out-of-order RISC-V CPU cores. He is also "CHERI in SoC" Work Group Lead at the CHERI Alliance, where his current work focuses on the practical system-level adoption of CHERI technology, including how CHERI-aware processors, interconnects, memories, and peripherals can be integrated into interoperable SoCs. Ben has more than 25 years of experience in semiconductor and SoC development, spanning architecture, design, verification, and software. Before joining Codasip, he held engineering leadership roles at Sondrel, an engineering management role at Imagination Technologies, and technical roles at Broadcom and Oxford Semiconductor.

References

CHERI-Zephyr and Compartmentalisation of Userspace Threads

J. Jackson
University of Birmingham, Birmingham, United
Kingdom



UNIVERSITY OF
BIRMINGHAM

Abstract

This talk demonstrates how isolated CHERI compartments can be built on top of Zephyr's existing userspace thread model, showing that PMP-based isolation can be replaced by hardware-enforced capabilities while granting software components only the authority they require.

Zephyr's userspace threads are a natural fit for building high-level software compartments because they already provide many of the isolation properties that compartments require. Today, this isolation is enforced through kernel-managed permissions and architecture-specific memory protection mechanisms such as RISC-V's PMP (Physical Memory Protection). CHERI enables the same security guarantees to be enforced directly through hardware capabilities, providing fine-grained, bounded access to memory without the region limits and domain-switching overheads associated with PMP. For example, each userspace thread is allocated its own stack, with isolation currently enforced through PMP. Under CHERI, this boundary can instead be enforced through a bounded stack capability installed by the kernel on entry, ensuring that out-of-bounds accesses are detected as hardware capability violations. Likewise, existing userspace system call interfaces map naturally onto CHERI compartment entry points, enabling controlled transfers of execution and authority between isolated software components. This enables developers to fully compartmentalise the library code their applications rely on to mitigate exploits or supply chain attacks. During the talk we will introduce CHERI-Zephyr and its development, explore how Zephyr's userspace model aligns with CHERI compartments, and then present an AES encryption library case study implemented as an isolated proof-of-concept compartment.

Biography

Jennifer Jackson

Jennifer Jackson is a Research Fellow at the University of Birmingham, where she leads the development of CHERI-Zephyr and the CHERI-Zephyr Working Group. She has worked extensively with CHERI technologies since 2021, including Arm Morello and CHERI-RISC-V platforms, with a focus on applying capability-based security to embedded and real-time systems. With a background in Electronic Engineering, she brings over 25 years of engineering and research experience spanning industry and academia. Her interests include memory safety, software compartmentalisation, computer architecture, and the co-development of hardware, firmware, and software, alongside FPGA-based system design and collaborative Git-based open-source development.

Jesse Spielman

Jesse Spielman is a postdoctoral researcher at the University of Durham. He works on the CHERI-Zephyr project implementing compartmentalisation features using CHERI technology. He recently completed his PhD at the University of Birmingham on the topic of using side channels to attack neural network activation functions. He is now developing expertise in the CHERI ecosystem and cybersecurity more broadly, having previously worked for 10 years on the technical side of film and TV visual effects production. He enjoys figuring out why things aren't working.

References

Compartmentalizing a CANopen stack with CHERIoT

H. Melder

Technical University of Munich, Munich, Germany

Abstract

Based on the CHERIoT architecture, we propose two architectures to restrict the blast radius of a potential exploit in an embedded system with a CANopen protocol stack. In both architectures, software components are compartmentalized and CAN frames filtered. The first architecture provides a lightweight firewall designed for reusing legacy C/C++ CAN stacks. The second architecture also features a firewall, but additionally offers strict flow isolation via sockets.

We implement these architectures on a CHERIoT evaluation board. The resilience of the system is demonstrated through fault injection, showing safe recovery of crashed components without compromising the entire system. Its real-time viability is assessed by measuring latency and compartmentalization overhead.

Biography

Hugo Melder is a computer scientist interested in compilers, programming language runtimes, embedded systems, hardware acceleration, and security research.

Over the past four years, he has worked across both hardware and software, including porting the Objective-C runtime to new architectures, designing FPGA accelerators, and contributing to open-source projects such as LLVM, GNUstep, and Apple's Grand Central Dispatch. His work also includes reverse engineering and security research.

Hugo recently completed a Bachelor of Science in Informatics at the Technical University of Munich with a minor in Electrical Engineering. He is also a co-founder of FAFO e.V., a nonprofit focused on semiconductor research and applied science.

References

Efficient Linkage-Based Compartmentalization on CHERI

D. Gao
University of Cambridge, Cambridge, United
Kingdom



Abstract

We present an efficient linkage-based model for in-process compartmentalization built on CHERI memory safety.

Our techniques enable the fine-grained compartmentalization of the entire UNIX user-space, which can scale to 10K+ compartments on desktop systems.

The model's "push-button" compartmentalization along existing library boundaries regularly hosts 500+ compartments per process for large applications such as Chromium, far exceeding the number of concurrently available protection domains supported by other mechanisms (e.g., up to 16 for Intel MPK).

Furthermore, custom policies can divide libraries into even more granular compartments.

Of the thousands of C/C++ programs we have tested, only the V8 JavaScript engine required source-level adaptation (<300 lines of changed code concerning garbage collection and JIT compilation) to operate with compartmentalization.

We implement the model for CHERI-extended versions of Armv8-A and RISC-V through support in the compiler toolchain and operating system, demonstrating the benefits of a single-address-space model, such as the smooth delegation of memory between compartments, as well as compartment-aware debugging and visualization.

We evaluate using multiple processors, including Arm's superscalar Morello and, notably, the first commercial CHERI-enabled RISC-V application core—Codalip's in-order dual-issue X730.

Biography

Dapeng Gao is working towards a PhD on CHERI compartmentalization at the University of Cambridge under the supervision Prof. Robert N. M. Watson. Previously, he worked on the formal verification of a CHERI-enabled RISC-V processor for his master's thesis.

References

LittleCHERI: Ultra-Fine-Grain Compartmentalization using Secure Stack-Sharing on CHERI Capability Machines

E. Storme
KU Leuven, DistriNet, Leuven, Belgium

The logo of KU Leuven, featuring the text "KU LEUVEN" in white capital letters on a blue rectangular background.

Abstract

CHERI capability machines provide good support for low-overhead enforcement of spatial memory safety and compartment isolation. However, compartment granularity and temporal stack safety remain constrained by the use of per-compartment stacks, a trusted intermediary switcher and/or compressed representability of precise stack capabilities. Theoretical support for secure stack sharing between distrusting compartments is strong but ignores practical concerns, and hence has not been implemented or evaluated.

We contribute LittleCHERI: a prototype secure calling convention that makes ultra-fine-grain compartmentalization viable on CHERI-RISC-V through secure stack sharing. LittleCHERI employs recently-proposed uninitialized and indirect sentry capabilities and contributes particularly a novel Reserve Stack technique that reconciles precise stack capabilities with capability compression. Benchmarks of our Clang/LLVM implementation show that LittleCHERI offers realistic overheads wrt. execution speed, stack pressure, and binary size, demonstrating that ultra-fine-grain compartmentalization is feasible on general capability machines.

Biography

Elias Storme is a PhD researcher at KU Leuven, working under the supervision of Dominique Devriese. His research focuses on the practical aspects of secure calling conventions, exploring the gaps between formal security results and real-world applications. Passionate about systems research, Elias aims to put novel security techniques into practice.

References

SignetOS: Designing a Secure and Fast Single-Address-Space Microkernel to Support Modern Workloads

J. Kressel

Google, Mountain View, United States of America

Abstract

Most modern operating systems task the memory management unit (MMU) with dual responsibilities: managing virtual-to-physical memory translation and enforcing isolation. Relying on page tables for security incurs steep costs: fragmented address spaces, TLB pressure, expensive context switches, and complex inter-process communication (IPC). While CHERI capabilities can replace the MMU for robust, hardware-enforced isolation, existing systems are largely restricted to statically configured embedded domains or struggle to support rich, dynamic, general-purpose workloads within a single address space.

In this talk, we present SignetOS: a CHERI microkernel that decouples memory translation from protection within a single 64-bit virtual address space and a single CPU privilege level. Rather than managing isolated address spaces, SignetOS retains the MMU solely for paging, enforcing all spatial and temporal boundaries via capabilities. Central to this design is the orthogonal decoupling of protection domains (compartments) from execution contexts (threads). Threads migrate across compartments via a lightweight domain switcher, reducing cross-domain security transitions to near-function-call overhead while eliminating page table switches, TLB flushes and message copying. To manage resources without complex resource managers, SignetOS governs access to system resources such as memory and CPU time through hierarchical quotas which can be used or sub-divided and delegated by the holder. SignetOS maintains a minimal, verifiable microkernel.

SignetOS also enables efficient memory-sharing optimizations for modern template-based application cloning. Mainstream platforms like Android rely heavily on `fork()` and Copy-on-Write (CoW) via a zygote process to accelerate application startup, an optimization that conventionally depends on duplicate virtual address spaces so absolute pointers remain valid across clones. SignetOS brings these memory-sharing optimizations to a SASOS through a "fork-less" CoW primitive paired with relative addressing meaning that pre-warmed runtime templates can be cloned into new compartments with on-demand physical page sharing and no pointer fixups.

We will present our system architecture, security guarantees, and current progress on SignetOS, demonstrating that CHERI capability-enforced security combined with a single address space offers a viable, highly performant foundation for modern workloads.

Biography

John Alistair Kressel

John Alistair Kressel is a student researcher at Google, researching single-address-space OSes and security with a particular focus on building secure and performant systems using CHERI. He is pursuing a PhD from the University of Manchester with a focus on the same topic.

Murali Vijayaraghavan

Formal Methods and Security Researcher at Google.

References

Ftrace on CHERI Linux - porting a fairly complex kernel subsystem

M. Kaiser
Codasip, Munich, Germany



Abstract

Codasip have ported the Linux kernel's tracing subsystem to CHERI. This talk provides an overview of the current status and describes some of the challenges we met along the way.

ftrace is another example where software could be ported to cheri with reasonable effort. The CHERI port uncovered several bugs in the existing codebase.

This session aims to raise awareness of the ongoing work and engage the audience for feedback, especially about missing functionality and potential extensions of tracing for CHERI.

Biography

Martin Kaiser is a senior software engineer at Codasip.

As a member of the operating systems team, he works on porting the Linux Kernel to CHERI. He has also contributed to the CHERI version of Qemu.

References

Title withheld (paper under submission)

N. Jordan
Ericsson Security Research, Stockholm, Sweden

Abstract

Abstract withheld due to the paper being under submission.

Biography

Nils Jordan is a MSc Cybersecurity student at KTH Royal Institute of Technology in Stockholm, Sweden, where he focuses on cybersecurity and especially software and systems security. He holds a BSc in IT-Security, providing him with a strong background in secure software and system design. During his internship at Ericsson Security Research, he contributed to developing an approach for complete and scalable temporal safety for CHERI application processors. This work introduced him to memory safety research and approaches for achieving memory safety through both software- and hardware-based solutions. His background in cybersecurity, combined with his practical research experience with CHERI, provides a solid foundation for investigating security challenges related to memory safety.

References

Using existing CHERI semantics for CHERI-native temporal safety

Y. Wang
University of Cambridge, Cambridge, United
Kingdom

Abstract

Unlike CHERI's native support for spatial memory safety, CHERI lacks inherent architectural support for temporal memory safety: capabilities can be freely copied and may remain reachable after the lifetime of the allocation they reference has ended. Existing approaches typically require additional metadata, memory scanning, or substantial software support, introducing performance and memory overheads. This talk explores how architectural support for temporal safety can instead be integrated with CHERI while avoiding external metadata and shadow tables.

We will present the key design choices developed across three recent works: using capability bounds to protect temporal metadata within allocations, encoding temporal state directly into capabilities, and exploiting CHERI's hierarchical capability semantics to manage temporal state efficiently. We will then show how these ideas can be extended to support repeated reuse of the same memory. Together, these design choices demonstrate how existing CHERI semantics can be leveraged to provide efficient temporal safety with broad security guarantees. These ideas have resulted in three work of mine so far: two have been accepted at top-tier security conferences, while the other one is currently under review.

Biography

I'm a second year PhD student, supervised by Prof Simon Moore. My PhD is about explore architectural extension to CHERI that supports efficient temporal memory safety. During the second-year of my PhD, I have successfully produced three paper on CHERI temporal safety, two have been accepted by top-tier security conferences, one is current under review by another top-tier conference.

References

Valid, In-Bounds, and Wrong: Analysis of Fault Injection Attacks Against CHERI Capabilities

J. P. Thoma
Codasip, Munich, Germany



Abstract

Memory-safety vulnerabilities remain the root cause of a large share of real-world exploits, and CHERI (Capability Hardware Enhanced RISC Instructions) has become one of the most promising hardware answers, attracting significant interest from academia and industry. Against the software attacker it was designed for, CHERI's capability mechanism works nicely. Deployed devices, however, are frequently physically accessible, and a physical attacker can inject faults that go well beyond the fault model assumed during CHERI's development. Memory safety and fault resistance are orthogonal properties, which raises the question of whether CHERI's guarantees survive a fault attacker at all.

We present the first systematic fault analysis of CHERI on RISC-V systems. We examine how instruction skips and bit flips affect capabilities and capability management instructions, and our threat analysis identifies which parts of the capability encoding are attractive targets and where the architecture already constrains the attacker. To evaluate generic CHERI-RISC-V systems, we built a functional fault simulator on top of Spike. With it, we demonstrate faults that violate CHERI's memory safety guarantees in simulation, and we assess countermeasures ranging from error-detecting capability encodings to fault-resistant control flow, including what each costs an implementer.

Biography

Jan Philipp Thoma

Jan Philipp Thoma is a security engineer at Codasip, where he works on CHERI-enabled RISC-V processor IP. His role includes threat modeling, supporting the architecture definition, and microarchitectural security analysis, with a current focus on the security properties of out-of-order cores that support capability-based memory safety.

He received his doctorate from Ruhr University Bochum in 2024 under the supervision of Tim Güneysu, with a dissertation on hardware designs for secure microarchitectures. His research sits at the intersection of capability security, microarchitectural attacks, and hardware fault injection. His published work centers on cache and TLB side channels, transient execution attacks, and on architectural defenses against them.

Maja Malenko

Maja Malenko works on CHERI RISC-V and microarchitectural security at Codasip. She is a CHERI Alliance Ambassador, and previously researched fault injection, side-channel attacks, and fine-grained memory isolation for embedded systems.

Yan Yan

Yan Yan works on CHE security at Codasip. Before joining Codasip, he worked as a researcher on cryptographic implementations focusing on side channel analysis.

References

Intra-allocation bounds enforcement in CheriBSD

A. Mazzinghi
University of Cambridge, Cambridge, United
Kingdom



Abstract

CheriBSD is the first UNIX operating system kernel with mature CHERI spatial memory safety support.

CheriBSD kernel allocators introduce bounds enforcement guarantees as part of the API contract; however, there are additional opportunities to narrow capability bounds to minimise software privileges. Richardson (UCAM-CL-TR-949) introduces compiler-assisted sub-object bounds to automatically narrow capability bounds with the C address-of operator, so that pointers to structure members are automatically bound to prevent access to adjacent members. Importantly, this also applies to array members, which are more likely to be indexed with a dynamically calculated value. This technique is valuable to mitigate memory corruption attacks that only partially overwrite the contents of a structure, without overflowing the allocation boundaries. In practice, we have found this to be relevant for a number of past FreeBSD kernel vulnerabilities and has real impact on the mitigation rate achieved by CHERI memory safety. As part of our work, the CheriBSD kernel has become the largest body of software shipping with compiler-assisted sub-object bounds enabled by default.

While sub-object bounds are valuable, they are currently considered to be a best-effort feature. This means that bounds may not precisely constrain access to nearby structure members under certain alignment conditions. Furthermore, sub-object bounds only apply to those cases in which the compiler can understand that the code derives a pointer to a sub-object. This talk introduces the problem of intra-allocation bounds enforcement in CHERI, grounding the discussion with quantitative and qualitative evaluation data from the CheriBSD implementation. While this discussion uses CheriBSD as a reference implementation, it applies equally to other operating system kernels, such as Linux, and user-space software.

As part of our work on CheriBSD, we identify a wider class of intra-allocation bounds use cases and evaluate implications for further bounds enforcement opportunities.

We identify three classes of intra-allocation bounds enforcement: compiler-assisted sub-object bounds, sub-allocation code patterns in kernel C code, and domain-specific bounds enforcement. In all these cases, a region of address space is further split into different logical components, possibly with different data types, which are accessed separately via pointer indirection. While the compiler toolchain can automatically handle most sub-object bounds enforcement, the other cases are not straightforward to detect and resolve. We contribute different approaches, developed within CheriBSD, to properly narrow bounds for these additional intra-allocation categories.

At the same time, we introduce the problem of bounds representability for intra-allocation bounds. CHERI capability compression schemes introduce additional alignment requirements for large objects. The same restrictions apply to pointers to different regions within a larger address space allocation. In order to ensure precise intra-allocation separation, each separate region must be constructed taking into account these alignment rules. This has a significant impact on the security guarantees offered by intra-allocation bounds enforcement. We show how this affects the different classes of intra-allocation bounds; in particular, we measure the existence of unrepresentable sub-object bounds in the CheriBSD kernel and discuss how the issue has been addressed. This is especially relevant for large arrays embedded in kernel structures and variable-length arrays.

Ultimately, we evaluate practical strategies to resolve these challenges, showcasing source code refactoring techniques, toolchain enhancements, and fail-close semantics that maximize the security guarantees of intra-allocation bounds.

Biography

Alfredo is a Senior Research Engineer at University of Cambridge and Research Engineer at Capabilities Limited.

Their main research interest is on memory safety in operating systems, but also in temporal safety, tracing and program analysis. They have developed the initial CHERI pure-capability port of the CheriBSD operating system, which is currently one of the most mature CHERI-enabled code-bases.

Other notable interests are the CHERI RISC-V standardisation effort, as well as the CHERI extensions to QEMU. They have also contributed to ports of various other software to CHERI, such as nginx, gRPC, CPython and Redis.

References

Writing applications for CHERIoT in Rust

E. Marangoni
SCI Semiconductor, Sheffield, United Kingdom



Abstract

The project to make Rust work on the CHERIoT platform has been ongoing for a bit more than one year. In this year we have achieved many results, starting from compiling parts of the standard library to CHERIoT. While this initial work is fundamental, writing firmware for the CHERIoT platform requires specific features to be exposed to users: for example, how can one derive a valid capability that points to an MMIO-bound area of memory in pure Rust? How can one write a compartment in pure Rust? What about a shared library?

We are now approaching a state where all the basic features to write firmware for CHERIoT in pure Rust are available, and the objective of this talk is to showcase these new features, and how they answer the questions above. Some of these features are CHERIoT-specific, while others are adaptable to every other CHERI platform. To this end, we will show multiple examples, each one explaining how to use the feature in case and how it works under the hood.

The first example will show how to derive valid MMIO-bound capabilities using the ``cheriot_mmio`` attribute: we will show how to write a driver for the LCD on the ICENI MPW1 board in pure Rust. The second example will show how to declare and define statically-sealed values in Rust with the ``cheriot_sealed`` attribute, writing an interface for the CHERIoT-RTOS allocator in Rust. This example will also show how to use the API to dynamically allocate new sealed capabilities. The third and final example will show how to declare compartments in Rust with the ``cheri_compartment`` attribute and how to handle errors in Rust. This example will also show how to design security boundaries with compartments working on real-world CVEs that have affected Rust crates.

Biography

Edoardo has been working on compilers for the majority of his professional career. After working in the Wasm space, he joined SCI Semiconductor and started working on the Rust-on-CHERIoT project since its beginning. He has been involved with every aspect of it, ranging from technical tasks such as adding the CHERIoT target to organising work for the team and communicating about the project to make it more visible.

References