IoT Security Foundation

JANUARY 2025

GLOBAL MARKETING STRATEGY REPORT

FOR SNBD-CHERI BASED ROUTER AND CONNECTED DEVICE INFRASTRUCTURE

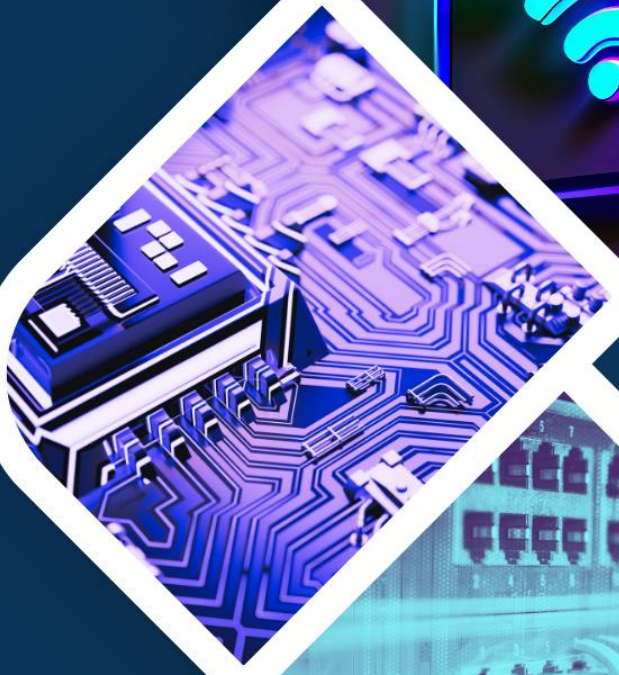

Lead Author

Haydn Povey

CEO of SCI Semiconductor & founding partner of the CHERI Alliance.

Edited by

John Moor

Managing Director of the IoT Security Foundation.

Acknowledgements

In the technology industry we all stand on the shoulders of giants. This report is no different and builds on substantial advancements across the industry around secure networking and memory safe computing. The industry, at many levels, moves incredibly quickly with competition, collaboration and individual delivery challenging us to do more, to implement secure solutions more quickly, and to resolve many of the most pernicious industrial challenges.

Significant thanks is given to the IoT Security Foundation (a membership service of Techworks) for enabling this report. Whilst the report is in support of the Secure Networking by Design (SNbD) collaborative project[1], it also highlights the importance of identifying and prioritizing the tasks to drive Memory Safe technology. Further thanks are given to all of the contributors into this report, with significant time being given freely to drive knowledge and adoption, to question the current status, and to ask the hard questions.

About the Lead Author

Haydn Povey is a technology and product cyber-security industry veteran, having helped to form the IoT Security Foundation, and serving on its Executive Steering Board since formation. Having led the successful launch of the initial Arm deeply-embedded and microcontroller product family (Cortex-M) he subsequently led the Arm security technology portfolio, including TrustZone and SecurCore, working with leading OEMs globally. He subsequently has led startups, including Secure Thingz (acquired by IAR Systems) and is currently CEO of SCI Semiconductor, a founding partner of the CHERI Alliance.

Editor

John Moor, IoT Security Foundation

Reviewers

Nick Allott, nquiring**minds**
James Willison, IoT Security Foundation

[1] https://manysecured.net/snbd/

Executive Summary

Given the reliance on computing in every aspect of society, it is no wonder that cyber security is one of the most critical global concerns for governments. McKinsey & Co recently estimated the impact of cyber-attacks on the global economy at $10.5 Trillion annually, equating to just under 10% of the worldwide GDP. Whether this is precisely correct, or not, the impact is certainly in the Trillions.

The importance of cybersecurity in an increasingly digital world was underlined in late 2024 by the Salt Typhoon attack[2]. The extent of this pernicious attack has yet to be fully realised, but certainly many tens of thousands of critical telecom systems have been compromised, leaking privileged governmental and commercial discussions, and enabling foreign entities.

As exposed yet again by Salt Typhoon, existing cyber-security methodologies are failing to keep pace with this malevolent problem, and the harsh reality is that the global industry must step up and take significant responsibility for this. The nature of complex systems and the capitalist society we live in demands products faster, lower cost, and with more features, with fierce competition that destroys companies in months if they fail to keep up. **The fundamental challenge is one of software, and specifically the industrial leveraging of open-source and third-party software that organisations integrate with little knowledge of contents, insufficient exploration and testing, and no regard for the critical vulnerabilities and malicious code that may be lurking in modern repositories.** In most cases, open-source code is generated by legions of smart, loyal, and dedicated engineers, but the harsh reality is that it is open to significant abuse, and the nature of complexity will always leave gaps open to exploitation.

The challenges highlighted by Salt Typhoon and a wide array of other attacks are defined by Memory Safety. Rather than being physically attacked, memory safety technically defines the interaction of software components on a digital device. As has been highlighted many times since the 1970's, computers are designed to be permissive systems, and this leads to a wide set of challenges, including memory overflows that over-write code with attacks, and pointer escalation attacks where the frameworks used to move around code are misappropriated, leading to malevolent attacks. For the past 50 years, the industry has attempted to manage this problem by using formal design methods (specify, implement, test), and more advanced memory management units. Ultimately these have failed – we must fix the foundations.

[2] https://www.cybersecuritydive.com/news/telecom-hack-salt-typhoon-china/734686/

More modern systems look to utilise memory-safe languages, such as Rust, to attempt to solve these issues, but again these are limited by the skills required to use them, the developer base of engineers, and the behemoth task of translating billions of lines of code from C/C++. While Rust and other advanced languages are very useful tools, the reality is software will always let us down because the people who write it are human - even the modern LLM (AI) systems are trained on sub-par code. The solution, as identified by many academic and governmental organisations, is a new hardware enforcement technology, namely **CHERI**.

CHERI, or Capability Hardware Enabled RISC Instructions, is an extension to existing computing architectures that implement a set of rigorous limitations, ensuring the software **Principle of Intentionality** and the **Principle of Least Privilege**. Fundamentally, these principles ensure that the blocks of code can operate as intended, with strong isolation between compartments but robust sharing of information through a well-formed and guarded Application Programming Interface (API). The only way to transition compartments is to delegate trust from one compartment to another and then ensure the permissions are revoked when transitioning back. The effect of these capabilities is to ensure only memory within specific bounds can be accessed, and only with very clearly defined permissions – **they become memory-safe**. This effectively removes a wide array of critical vulnerabilities, with benchmarking demonstrating that 70% of critical vulnerabilities and exploits (CVEs) are prevented.

The ManySecured[3] program, managed by the IoT Security Foundation, and the Secure Networking by Design (SNbD) project which flowed from it, have formed a beachhead demonstrator for the CHERI technology. As seen by Salt Typhoon, **telecommunications and networking are primary targets of nation-state attacks, both because of the widespread repeatability of the attacks, and the critical information that flows over them**. Even when encrypted the harvesting of critical information for later analysis and decryption is rife at the nation-state level, leaving companies and governments critically exposed.

Leveraging the Arm Morello test chip - an early implementation of the CHERI technology available as an Arm Neoverse high-performance processor - the SNbD project has clearly demonstrated the potential impact of CHERI, providing a framework for future projects and products. This report exposes more details of this audacious program; however, fundamentally it has been proven that we are standing on the verge of a new paradigm in secured computing. The ability to develop memory-safe and

[3] https://manysecured.net

compartmentalised applications that are inherently self-sealing and immune from such a wide variety of modern attacks represents a major milestone in the industry.

Beyond the SNbD project's route to market analysis, this report also justifies a redoubling of the current efforts in securing the digital landscape. The foundation technology is now proven, and it is up to the industry, governments and other critical stakeholders to drive it to its fruition. To avoid future pervasive nation-state attacks the industry must be incentivised to adopt the new fundamental components of CHERI-enabled silicon, compartmentalised and memory-safe operating systems, and updated tools including CHERI-aware compilers and memory-safe languages such as Rust. As the ManySecured Secure Networking by Design program has shown, we are within touching distance and should complete our transition to a secure-by-design future.

Contents

Introduction

In October 2022 McKinsey & Company released a highly significant report on their survey of cybersecurity practices[4] outlining the massive challenges the world faces from cyber-attacks, confidentiality breaches, nation-state competition and general malfeasance. The survey suggests that the global impact of cyberattacks will continue to rise, and by 2025 will reach $10.5 Trillion annually, a growth of 300% from 2015 measurements. This huge number represents a significant percentage (9%) of the annual global GDP of approximately $117 Trillion, and at one level seems incongruous. Yet viewing the widespread impact of ransomware attacks on hospital services, widespread attacks on the general population to steal personal data and banking details, and industrial-scale theft of intellectual property, seems to pass the collective sniff test. Beyond nation-state warfare, cybersecurity must now rank as the largest threat to global economies across all regions, and no single country or industry is immune to the threat.

With the global economy so interdependent and the digitalisation of platforms continuing to accelerate, the security of our Information Technology (IT) and Operational Technology (OT) becomes the leading battlefront for organisations and nation-states. The cyber defence IT & OT marketplace is estimated at $250 Billion per annum by McKinsey, but again, the report paints a very dark picture of the reality, with this potentially growing to $2 Trillion per annum to meet the evolving challenges from our enemies. These numbers are astounding, especially as it can be seen that, while we have many successes, overall, the industry is struggling to meet the challenge. In part this is the human aspects, with phishing attacks rising; in part it is the rise of AI, driving more sophisticated and pernicious attacks; but largely, it is because the fundamental computing technology we all rely on was never intended to meet this level of threat, with a permissive, rather than restrictive, architecture. This challenge is reinforced through the nature of competitive business and the race to get products to market at the lowest cost possible, winning critical market share and ultimately boosting profits. These natural requirements drive up code reuse unsustainably while limiting the amount of testing which can be completed. While practices such as Continuous Integration and Continuous Deployment (CI/CD) help in closing the testing gap, the sheer amount of 3rd party code that projects inherit means substantial zero-day attacks will propagate long into the future; that well-placed malignant code will continue to exist in systems; and that simple short-cuts that are appropriate for one project will deliver critical flaws in others.

[4] https://www.mckinsey.com/capabilities/risk-and-resilience/our-insights/cybersecurity/new-survey-reveals-2-trillion-dollar-market-opportunity-for-cybersecurity-technology-and-service-providers#/

Purpose

The purpose of this report is to highlight the progress made to resolve one of the most pernicious issues of our time, memory safety, and to highlight the gaps remaining to convert a promising cybersecurity project into global change.

The report is structured to review the challenges we face in the general digital domain, and explicitly the communications and networking market, through the prism of the router, traditionally at the sharp end of the network infrastructure, where security pressures most tightly compete with cost, usability, and a fast-moving environment. The Secure Networking by Design (SNbD) project was envisioned as a mechanism to identify critical flaws in modern systems and the best routes to resolve them based on CHERI technology. The project has been successful in achieving this, and this report is first and foremost a written record of the outcomes of that project.

This report's primary focus is addressing the key challenge: bringing "digitally secure by design products" to market. It analyses this issue and proposes solutions in the accompanying annex. There are three major aspects to this, which are dealt with explicitly.

Firstly, there are inherent systemic issues in the way our digital landscape is constructed today. Primarily based on very well-known memory safety issues, that were first identified over 50 years ago, the reality is that neither industries nor governments have been able to resolve these… until now.

Second, there are explicit challenges in resolving these issues through the availability of technology, stretching from hardware availability to tooling and memory-safe languages, and ultimately the skills and knowledge pool that is known to be small. There are very few programmers who create vulnerable-free code, and there are many who leverage vulnerable codebases, thus we have to solve the problem ‘for the many’.

Third, there is a major gap evolving between the requirements the government is putting on industry, and the industry’s knowledge and capabilities to resolve these demands. In the USA, CISA and the FBI recently requested that all OEMs and device suppliers resolve a roadmap to memory-safe systems by the end of 2025. However, it is clear that the vast majority of organisations neither have the knowledge nor desire to meet these challenges. Collectively, we have a massive task of education and enablement, and this report, alongside the SNbD project, is looking to resolve this.

Methodology

This report is based on a standard evaluation methodology to identify critical next steps for Secure Networking by Design activities.

Information has been sought across a wide array of technical resources covering:

- Upcoming cyber threats and emerging cyber landscape reports
- Governmental guidance from the Office for National Cyber Director at the US White House, and others
- Regulatory best practices, including EN303645, UK's PSTI, EU's CRA and others
- Technical analysis of routers and networking requirements
- Memory Safety technical documentation

Additionally interviews and conversations have been undertaken with significant stakeholders covering

IP vendors
Silicon device vendors
Memory Safe technology stakeholders
Networking device manufacturers
IoT device manufacturers
Governmental stakeholders (UK/US)

SECTION 1: Memory Safety

The Importance of Memory-Safe Operation

As outlined in the introduction, the issues around cybersecurity continue to grow to the point at which global industry and infrastructure will soon be overwhelmed. The sheer complexity of modern software stacks, leveraging commercial and open-source 3rd party code, to solve incredibly multifaceted issues, means we will always be faced with making a simple choice between getting code correct or getting products shipped. It is simply unviable to test every code fragment against every potential exploit, and while rewriting some critical code components in advanced Memory Safe languages, such as RUST, it is also simply unaffordable for the industry to recreate large applications, operating systems, or even embedded solutions.

The computing industry has recognised these challenges for many decades, with memory safety reported as an issue back in 1972 and beyond. The advent of mass connectivity and the modern Internet took these issues from theoretical to implementable, with the Morris Worm[5] implementation in 1988, potentially the first Internet worm that exploited simple buffer overflows. Computer architectures have subsequently attempted to mitigate these flaws through the advent of simple Memory Protection Units and Memory Management Units; limited entry points and multi-layered network stacks; and secure or trusted execution environments to reduce the attack surface. However, as the architectures are permissive and the code environment very large, these have all been breached in the real world.

Major steps have been taken in domains where network and computer failures are unacceptable, such as military and aerospace systems. In 1975 the US Department of Defence, concerned by the lack of safe modular programming, created a High Order Language Working Group to review these needs, and subsequently, this working group evolved the Ada programming language. This language was strongly accepted at the time and was widely believed to become the dominant programming language, but due to challenges in code size, complexity, and cost, has languished against lighter, more flexible and more popular languages, today ranking #98 on the list of the Top 100 Programming Languages[6] (by social mentions).

Government Recognition of Memory Safety in Cybersecurity

Today we see a renewed focus on Memory Safety, with the publication in December 2023 of a call for action from international cybersecurity authorities, including US agencies (CISA, NSA, FBI) and the governments of Australia, Canada New Zealand and the UK. The guide "*The Case for Memory Safe Roadmaps: Why Both C-Suite Executives*

[5] https://en.wikipedia.org/wiki/Morris_worm
[6] https://www.libhunt.com/index (June 2024)

and Technical Experts Need to Take Memory Safe Coding Seriously"[7], is a well-constructed paper aimed at making the issues simpler to understand for stakeholders of all levels within organisations and companies, with several call-to-actions. However, unless the technology is available, there and then, it is challenging for businesses to alter course.

As often highlighted, memory safety has been identified as the root cause for over 70% of the CVE's in systems running the Chromium browser engine[8]. Other systems will be impacted by differing amounts, but fundamentally, memory safety is the single biggest source of vulnerabilities in software. Addressing this issue has been a high priority for governments around the globe, both to prevent nation-state-sponsored attacks, and also to prevent the degradation of critical infrastructure and industries.

At the CyberUK'24 conference, both the CEO and CTO of the UK's National Cyber Security Centre, the public arm of GCHQ, highlighted in their keynote addresses that memory safety was their number 1 systemic issue and the need to urgently address it across the IT & OT estate.

The single biggest issue highlighted at this, and many other events is the cost to transition billions of lines of code to memory-safe languages. The reality is that this is simply too large a task and will never happen for cost and complexity reasons, alongside the challenge of knowing where to start when so much open-source code is embedded into platforms. While Rust is undoubtedly a very good memory-safe programming language, only new code and critical components of select systems will be authored here.

This leaves a huge legacy C /C++ issue which Rust will not address for strictly commercial reasons, which will continue to host issues for decades to come. As such, we are left with limited options and memory-safe hardware has strong appeal.

Mitigating Memory Safety Issues

It is important to clarify what is meant by Memory Safe issues, and which are resolved by Memory Safe languages and hardware. For example, there is often confusion between memory-safe ***execution of software*** where buffers are impacted, and unexpected behaviours are propagated, and ***memory safety***, or ***secure storage***, where memory blocks are protected against physical attack, or clocking out of code. In the context of this report, we are dealing with Memory Safe execution of software, and many of the explicit challenges outlined in the following table:

[7] https://www.cisa.gov/news-events/news/cisa-nsa-fbi-and-international-cybersecurity-authorities-publish-guide-case-memory-safe-roadmaps

[8] https://github.com/microsoft/MSRC-Security-Research/blob/master/papers/2020/Security%20analysis%20of%20CHERI%20ISA.pdf

Common Memory Safety Issues	
Access Error Attacks e.g. Invalid read / write of a pointer	
Buffer Overflow	A buffer overflow or overrun is an issue where a program writes data to a buffer beyond the buffer's allocated memory, overwriting adjacent memory locations. By writing into an overflown buffer, it is possible to inject malicious code into the system, subsequently impact system behaviour, leading information, and potentially taking malicious control.
Buffer Overread	A buffer overread is where a program reads from a buffer, but subsequently also reads adjacent memory. Mostly caused by misconfiguring boundary implementation this enables large portions of code to be accessed by an attacker who can then form highly targeted attacks.
Invalid page fault	A page fault, or hard fault, is where the system is forced to access a pointer outside of the virtual memory space often causing an exception which the attacker can use to gain access and corrupt operating systems. Typically targeting rich operating systems (e.g. Linux / Windows) and application processors.
Use After Free	Use after free is an attack vector to read the contents of memory after it has been used, and subsequently released (freed), without the contents being cleared. Often also known as a dangling pointer, the system normally eradicates these through a garbage collection mechanism, but valuable information may be available.
Uninitialized Variables Attacks e.g. a variable has been created but not value attributed at that point	
Null pointer	Dereferencing of a null, or unattributed, pointer enables an attacker to impact where in the memory this variable will now point. Given many pointers are created dynamically in the program execution impacting this value enables an attack free rein over impacting the memory system covering both code and data.
Wild pointers	A wild pointer is where a pointer is compromised and maliciously used before being initialised to a known state. This enables attackers to probe the system invisibly, unless the pointer state is checked prior to being set.
Memory Leaks Memory is tracked incorrectly enabled information to leak, build up or overfill in critical execution	
Stack Exhaustion	Poor code implementation, or a malicious actor may force the stack pointer to exceed the stack bounds, creating consequences of unknown form. Sometimes the systems may crash, spilling information into a publicly available state, or may create an unknown operation which an attacker can then impact.
Heap Exhaustion	Heap exhaustion is caused by the code attempting to allocate more memory that is physically available, causing an out of memory fault. Similar to stack exhaustion this can subsequently cause information to be inadvertently shared, or for maliciously operated control points to be evolved.
Double Free	Accidentally attempting to free memory multiple times may enable information to be shared, that should remain confidential, especially if

	additional execution is expected before the freeing of memory. For example, contents may still be held that would otherwise be blanked, or interim control values may be exposed.
Invalid Free	Attempting to free memory with an invalid address may have serious consequences, either in the attacker retargeting the free address to unlock critical control vectors, or where freeing of memory is carried out prematurely to get to the current contents. These attacks are often hard to constrain within systems if privilege levels have not been correctly set.
Mismatch Free	Similar to invalid free issues, but where multiple memory allocators may be in use, this potentially enables attackers to unlock instructions or data in critical sections, taking control, or highlighting additional attack points.
Unwanted Aliasing	If a memory is accessed through multiple aliases, either statically or dynamically set, there is opportunity to impact the result of one call by maliciously attacking the other. It is poor programming practice to implement multiple pointer calls to the same object, but is widely done as a shortcut instead of formally defining, and limiting, access to the object's location.

A complete list of Memory Safety vulnerabilities is identified on the Common Weakness Enumeration (CWE) website at https://cwe.mitre.org/data/definitions/1399.html which all developers are suggested to familiarise themselves with.

An indicative list is shown in Figure 1 below:

CWE CATEGORY: Comprehensive Categorization: Memory Safety

Category ID: 1399
Vulnerability Mapping: PROHIBITED

Summary

Weaknesses in this category are related to memory safety.

Membership

Nature	Type	ID	Name
MemberOf	V	1400	Comprehensive Categorization for Software Assurance Trends
HasMember	C	119	Improper Restriction of Operations within the Bounds of a Memory Buffer
HasMember	B	120	Buffer Copy without Checking Size of Input ('Classic Buffer Overflow')
HasMember	V	121	Stack-based Buffer Overflow
HasMember	V	122	Heap-based Buffer Overflow
HasMember	B	123	Write-what-where Condition
HasMember	B	124	Buffer Underwrite ('Buffer Underflow')
HasMember	B	125	Out-of-bounds Read
HasMember	V	126	Buffer Over-read
HasMember	V	127	Buffer Under-read
HasMember	V	129	Improper Validation of Array Index
HasMember	B	131	Incorrect Calculation of Buffer Size
HasMember	B	134	Use of Externally-Controlled Format String
HasMember	B	188	Reliance on Data/Memory Layout
HasMember	V	198	Use of Incorrect Byte Ordering
HasMember	V	244	Improper Clearing of Heap Memory Before Release ('Heap Inspection')
HasMember	V	401	Missing Release of Memory after Effective Lifetime
HasMember	V	415	Double Free
HasMember	V	416	Use After Free
HasMember	B	466	Return of Pointer Value Outside of Expected Range
HasMember	B	562	Return of Stack Variable Address
HasMember	V	587	Assignment of a Fixed Address to a Pointer
HasMember	V	590	Free of Memory not on the Heap
HasMember	∞	680	Integer Overflow to Buffer Overflow
HasMember	∞	690	Unchecked Return Value to NULL Pointer Dereference
HasMember	V	761	Free of Pointer not at Start of Buffer
HasMember	V	762	Mismatched Memory Management Routines
HasMember	B	763	Release of Invalid Pointer or Reference
HasMember	B	786	Access of Memory Location Before Start of Buffer
HasMember	B	787	Out-of-bounds Write
HasMember	B	788	Access of Memory Location After End of Buffer
HasMember	V	789	Memory Allocation with Excessive Size Value
HasMember	B	805	Buffer Access with Incorrect Length Value
HasMember	V	806	Buffer Access Using Size of Source Buffer
HasMember	B	822	Untrusted Pointer Dereference
HasMember	B	823	Use of Out-of-range Pointer Offset
HasMember	B	824	Access of Uninitialized Pointer
HasMember	B	825	Expired Pointer Dereference

Figure 1: CWE Memory Safety Vulnerabilities

A Path Toward Secure & Measurable Software & Systems

In February 2024 a significant intervention was made by the US White House Office of the National Cyber Director (ONCD) [9] in their report on "Back to the Building Blocks: A Path Toward Secure & Measurable Software"[10].

This report highlighted the need to secure the building blocks of cyberspace through three critical components, all of which are applicable to the development of Secure Networking by Design.

1. Memory Safe Programming Languages

The report requires that organisations writing code should leverage memory-safe languages and correctly states that multiple alternatives now exist. While being careful to remain apolitical and above the comparative language wars, it is clear that numerous choices are available here. The most notable is Rust, a multi-paradigm, general-purpose language, initially sponsored by Mozilla. Rust's syntax is similar to that of C and C++, although many of its features were influenced by functional programming languages such as OCaml[11]. It has been described as targeted at "frustrated C++ developers" while emphasizing features such as safety, control of memory layout, and concurrency. Rust is currently #6 on the top 100 Programming Language list, demonstrating excellent traction across the industry.

The report also highlights that using memory-safe programming languages for new products can provide significant advantages, but there are still at least three major challenges to address:

Challenge number 1; while RUST is currently growing very quickly in uptake and interest, it remains a specialised skill base. Typically, it has been identified that uptake is strongest in younger engineers who may not have the experience or training to bring it to bear in the industry. This is not always the case, and experience must not be diminished on age alone, but the harsh reality is that there are insufficient experienced engineers capable of leading the widespread adoption of the language

Challenge number 2; while RUST is as stated "targeted at frustrated C++ developers", the language may be less suitable for deeply embedded systems, including networking and routers, versus traditional C. There are a wide set of appliances in this domain, and hence many applications will be able to take advantage of RUST memory safety, but many will not.

Challenge number 3; is arguably the most crucial - the reality is that huge amounts of code already exist, and most will not be changed, reauthored, or ported. This is an unfortunate reality of modern product development, where most "new" products are

[9] https://www.whitehouse.gov/oncd/briefing-room/2024/02/26/press-release-technical-report/
[10] https://www.whitehouse.gov/wp-content/uploads/2024/02/Final-ONCD-Technical-Report.pdf
[11] https://en.wikipedia.org/wiki/Rust_(programming_language)

evolutionary, rather than revolutionary. It is far easier, cheaper, and commercially substantially low risk to build upon what has already been shipped, and hence while better solutions can be evolved this will only happen as legislation or regulation dictates. The DARPA TRACTOR program[12] is looking to assist in this process through the use of AI / LLM to translate code automatically, but as always the critical difficulties lay in the details of the code, the platform it is developed for, and the training data.

2. Memory Safe Hardware

The ONCD report specifically calls out the need for memory-safe hardware, with a clear desire for critical systems to be ***secure by design***. Explicitly calling out CHERI (Capability Hardware Enhanced RISC Instructions) in the report, it also discusses some alternative technologies, such as memory-tagging extensions (MTE). The MTE technology is inherently useful, especially around aspects of system debugging, yet is vulnerable to speculative execution attacks if used as a security barrier[13][14], against the developer's explicit advice.

3. Formal Methods

The final area called for improvement in critical systems, including telecommunications and networking, is that of formal methods, which are lacking in many of the popular systems deployed today. The lack of detailed system specification and explicit measurement and testing to the specification is a key weakness in many systems, introduced in the rush to get the product to market, with minimal engineering resources applied to drive down short-term costs. The consequences of this 'move fast and break things' approach are that a minimal viable product is often released to the market, with an intent to follow up with a more robust version 2.0 release. Unfortunately, the v2.0 release is never quite released due to market pressures and poorly designed products fester in our networks.

A robust call to action is represented in the ONCD paper, firstly for formal methods to be incorporated directly into the developer tools chain, and secondly for any third-party code to be formally verified before integration. While these again are highly laudable calls to action, the harsh reality is that organisations who want to produce quality products will invest in these development flows, a significant majority simply will not invest or do not have sufficient skills to utilise them, or do not understand the implications of a failure. In this matter, as with other software development flows, the foundational technology simply must pick up the challenge and ensure that bad habits can be supported and sustained without impacting developer efficiency. While better tools and simpler formal methods need to be evolved, ignoring the reality of how engineering operates today is what got the industry into the mess it is in today. As such we must develop hardware technology which mitigates the "blast radius" of vulnerabilities while supporting engineers to identify and fix poorly implemented code. The CHERI technology is, therefore, of significant merit.

[12] https://www.darpa.mil/program/translating-all-c-to-rust
[13] https://arxiv.org/abs/2406.08719
[14] https://developer.arm.com/documentation/109544/latest

Memory Safe Technology Deployment

The general adoption of memory-safe technology is regarded as a necessary step for all enterprises, and while much of the report focuses on SNbD and CHERI hardware, it is important to see this in a wider context. As such these three fundamental requirements have been identified by the White House Office of National Cyber Director and should be adopted into industry and government policy.

1. **Prioritisation guidance.**
   Manufacturers should be guided to consider how to prioritize migration to memory-safe technology, the near-term impact on product roadmaps and specific guidance for development and technical teams.

2. **Picking appropriate use cases for memory-safe technologies.**
   There are numerous approaches to memory-safe technology, and each one has its own set of trade-offs in terms of architecture, tooling, performance, popularity, cost, and other factors. While CHERI has substantial advantages no single approach is right for all programming needs. Manufacturers and system producers need to look at use cases individually and pick the most appropriate solution for each.

   When selecting an approach, software producers should follow standard risk management processes, as memory safety solutions are not free from other potential vulnerabilities of critical severity.

3. **Staff capabilities and resourcing.**
   Enterprises need to consider how they will train developers in their selected approach, how they can prioritize hiring developers with the relevant skills, and what resources they may need to support the selected technology. For example, with CHERI a recompile of existing C code is sufficient to resolve memory-safe issues on CHERI-enabled hardware, whereas for Rust a more thorough porting is required, impacting project resourcing and diverting critical resources.

SECTION 2: ManySecured and Secure Networking by Design

Introduction

The Secure Networking by Design (SNbD) project is driven by the ManySecured Working Group, a function of the IoT Security Foundation. It is a vehicle to address the needs of networking in alignment with a number of major stakeholders, including the UK government's Digital Security by Design program, to leverage aspects of the UKRI Technology Access Program and the Morello test platform. Significant progress has been achieved through this activity, which is outlined further below.

The ManySecured WG is "an open ecosystem, designed to improve network security against IoT attack through an **"intelligent defensive controller"**. The SNbD collaborative project has a clear goal of demonstrating how a future-generation CHERI-based secured hub can be used to manage constellations of devices which may become compromised. This hub provides real-time intelligence to monitor activity at the gateway, determine the threat level, and take appropriate action.

More information on ManySecured can be found at https://manysecured.net

ManySecured-SNbD Key Elements

Excellent progress to date has been made on ManySecured, based on the UKRI Technology Access Program and the Arm Morello test chip. This is a significant proof of concept, but sizeable obstacles to the adoption of CHERI-based systems are identified, and outlined further in this report, alongside suggested next steps.

The ManySecured project has themed subgroups as described here.

SNbD Workstream (Secure Networking by Design)	Incorporates secure networking by design concepts into the networking ecosystem, examining how the industry builds memory safe secure routers and networking devices using CHERI and other memory safe secure technologies. Expected outcomes include - Open-source implementation of a CHERI hardened router (code on Morello test chip) - Security analysis of the impacts of memory safe interventions (documentation) - Tools to evaluate impacts of memory safe interventions (code) https://specs.manysecured.net/snbd/

SUIB
(Secure Usable Internet Browser)
	The SUIB working group was formed to address the fundamental problem of how to securely connect a browser to a local (private network) web server, possibly hosted on an IoT device or router.
	
	Expected outcomes include:
	- Problem statement definition/whitepaper: a detailed overview of the problem scope
	- Technical requirements: high level technical requirements which:
	  ○ embody the problem statement
	  ○ forms a benchmark for evaluating the completeness & quality of proposed solutions
	
	https://specs.manysecured.net/suib/

D3
(Distributed Device Descriptors)
	The D3 workstream addresses two main challenges:
	
	1. How does a community make statements about device types (as opposed to device instances), reliably and securely?
	2. How can the community reason about devices reliably (human-centric or machine-centric)?
	
	D3 provides structured data of known provenance, which can be used to assert claims about how IoT devices should behave.
	
	Expected outcomes include:
	○ a fine-grained analysis of the problem scope
	○ high-level technical requirements which:
	  • embody the problem statement
	  • form a benchmark for evaluating the completeness & quality of a proposed solution
	○ a detailed technical document defining the formal interfaces and data schemas used to embody the solution
	
	https://specs.manysecured.net/d3/

D3Con (D3 Control)	The D3Con workstream addresses how to securely, and with interoperability, extract control of a router. In particular D3Con must be able to trigger actions that can act on individual devices or sets of devices to either protect devices or prevent devices doing further damage.

	Core behaviors are supported by DCon inducing: ○ Segmentation to constrain risks within classes of devices ○ Segmentation of Network ○ Allocation of devices to segments ○ Creation of inter-segment bridges and secure communication ○ Behavioural containment ○ Autonomous device disablement https://specs.manysecured.net/DCon/
D3Events	The D3Events workstream addresses how to securely, and with interoperability, extract security relevant metadata from gateway and router devices, in order to detect anomalous behaviour of IoT devices at scale Core behaviors supported by D3Events include • **Response**: provide sufficient real-time (or close to real-time) security relevant information about connected devices. • **Profiling**: provide sufficient historical data about connected devices • **Analysis**: provide sufficient historical data about all connected devices • **Forensics**: (optional) provide information to assist with forensic analysis of security events. Expected outcomes of this workstream include ○ Clear statement defining the problems to be solved ○ High level technical requirements ○ Detailed technical document defining the formal interfaces and data schemas to embody the solution https://specs.manysecured.net/D3Events/

Further working groups on the Router Threat Model (GCERT) and Lifecycle Management (Lifecycle) are also implemented and reflect excellent progress to date around a very challenging area.

SECTION 3: CHERI Technology, Scoping & Benefits

Many readers of this report may already be aware of CHERI as a technology and the various implementations which have already been produced. Please see Annex IV for a brief introduction if you are unfamiliar or require a refresh.

About CHERI

CHERI, or Capability Hardware Enhanced RISC Instructions, is a joint research project of SRI International and the University of Cambridge[15] to revisit fundamental design choices in hardware and software to dramatically improve system security.

CHERI extends conventional hardware Instruction-Set Architectures (ISAs) with new architectural features to enable fine-grained memory protection and highly scalable software compartmentalisation. The CHERI memory-protection features allow historically memory-unsafe programming languages such as C and C++ to be adapted to provide strong, compatible, and efficient protection against many currently widely exploited vulnerabilities. The CHERI scalable compartmentalisation features enable the fine-grained decomposition of operating system (OS) and application code, to limit the effects of security vulnerabilities in ways that are not supported by current architectures.

An Introduction to CHERI[16] white paper is available on the University of Cambridge website. It is an extremely well-written document and regarded as required reading to understand more about memory-safe computing, capabilities, compartmentalisation, and CHERI technology in general.

CHERI Functionality

It is important to emphasise some of the critical benefits, which are often lost when Memory Safety is mentioned.

The Two fundamental principles which are worthy of emphasis in this report are:

- Principle of least privilege
- Principle of intentionality

The principle of least privilege indicates that software objects should always be given just sufficient privileges to perform their tasks. The unfortunate reality is most programmers quickly escalate to supervisor or privileged mode and stay there, because it is simple, and they do not question what could go wrong. This is a fundamental flaw with third-party code, as most of it demands access it simply does not need.

[15] https://www.cl.cam.ac.uk/research/security/ctsrd/cheri/
[16] https://www.cl.cam.ac.uk/techreports/UCAM-CL-TR-941.pdf

The principle of intentionality is equally critical, and broadly misused in code today, where a simple call into a function is not bound to the calling function and therefore can be easily abused. Most notably pointers are easily abused, and many vulnerabilities can be traced back to not ensuring the code is used for the specific purpose intended. The intentionality of CHERI capabilities marks a hugely significant improvement in computing, albeit with limitations some programmers will dislike. The ability for a process to pass a capability as an argument to a system call, and subsequently enforce limits on how the function operates is a simple but powerful mechanism that every reader should be aware of.

Extracting CHERI Benefits

The challenge with a foundational technology, such as CHERI, is that you need to be an expert to understand the consequences of what it achieves, how it accomplishes it, and the benefits it delivers. In the world of memory safety and advanced computer science, that immediately disenfranchises 99.9999% of the population or more. As such it is essential that we talk in terms of the OEM developer and end-user benefits, to swing the commercial engagements that will form the bedrock of popularising the technology.

Ferocious Code Reuse

CHERI at its heart is designed to not shy away from the limitations of open-source and third-party code, but instead acknowledge the challenges and provide a prophylactic interface, where if malevolent code is present, it cannot infect the system. It is intended to enable ferocious code reuse and take significant limitations out of the development process.

A powerful advantage of CHERI - including CHERIoT, its smallest implementation - is quite simply that it enables code of varying quality to operate in a secure framework. That is, by using the foundational memory safe capabilities, and compartmentalisation, any critical software vulnerabilities that exist on the system can be trapped, constrained, and mitigated, enabling end-users to continue to use the product protected.

This realisable benefit cannot be overstated – it is incredibly important. Why? The reality of the technical world is few people, if any, create software from scratch. When given a task to program the first action of most developers is to search the web, GitHub, SourceForge, for example, and start torturing it into the desired share. This is not intended in any way to denigrate the task; it is just to recognise the reality. Where specific functions are required, such as a network stack, the programmer will look for a trusted 3rd party open-source vendor and program to the APIs. The challenge is that this code is itself complex, it will likely have flaws and may even host calls to malicious code that have specific backdoor functionality hidden away. The code will likely pass testing, even extended fuzzing, but may still have specific functionality built in that the user would not want to use if they could identify it.

The bleak reality is that the software industry is so addicted to “free” code that we could not move away from it if we wanted to. The commercial need to ship applications in minimum time and with minimal resources drives us to reuse code and ship.

Secure by Design: Isolation, Compartmentalisation Safety & Code Reuse

The principle of isolation is, in general, a critical one. The best way of protecting a system is to air-gap it and rip out any air interfaces, such as wifi. However, this being impractical in most modern systems has led to numerous alternatives, including memory protection units (MPU), memory management units (MMU), virtual machines and hypervisors, and so forth. The challenge with these process-bound solutions is typically not the isolation itself - this can be enforced robustly - but rather the isolation and subsequent sharing of information. These APIs can be made secure if subject to formal methods of design but are inherently overly permissive, with no explicit privileges. This is resolved in CHERI through the capabilities, replacing pointers with structures that both limit the addressable memory space and define a specific set of privileges for what the code can achieve. The capabilities further revoke the privilege status as the call unwinds, ensuring no open attack vectors remain.

The ability to create robust compartments is often overlooked but offers benefits of the same consequence as core memory safety technology. In a recent rebasing of the FreeRTOS Network stack to the CHERIoT processor system, it was shown that via memory safety plus compartmentalisation, the team were able to mitigate 100% of the CVEs registered on cve.org, with 7/10 being based on memory safety, and 3/10 mitigated through compartmentalisation.

CVE	Mitigated by CHEIRoT ISA properties	Mitigated by CHERIoT RTOS Properties	Mitigated by Compartment model
2018-16522	√	√	
2018-16526	√		
2018-16525	√		
2018-16599	√		
2018-16524	√		
2018-16527	√		
2018-16600	√		
2018-16602	√		
2018-16603	√		
2018-16523			√
2018-16598			√

Figure 2 FreeRTOS Network Stack CVE Resolution (Courtesy: SCI Semiconductor)

Of the three CVEs not being resolved via memory safety, it was demonstrated that these issues can be separated and bound to tightly limit any attack. The system prevents malignant escalation of privilege, and if a compartment is infected it can be rapidly resolved and reset by the remainder of the system. While not bulletproof, it does ensure

that if an attacker wishes to gain control, they need to have many exploits across several compartments, which becomes exponentially more challenging and therefore less attractive to them.

For developers, structuring the compartments is a significant task. While it is possible to move legacy code over into a simple, single compartment, thought should be taken as to how to build a robust and structured set of compartments, ensuring that a flaw in any one of them cannot drive an exploit, through management of the capabilities to limit the privilege state and ensure the intentional use of the functionality is encapsulated.

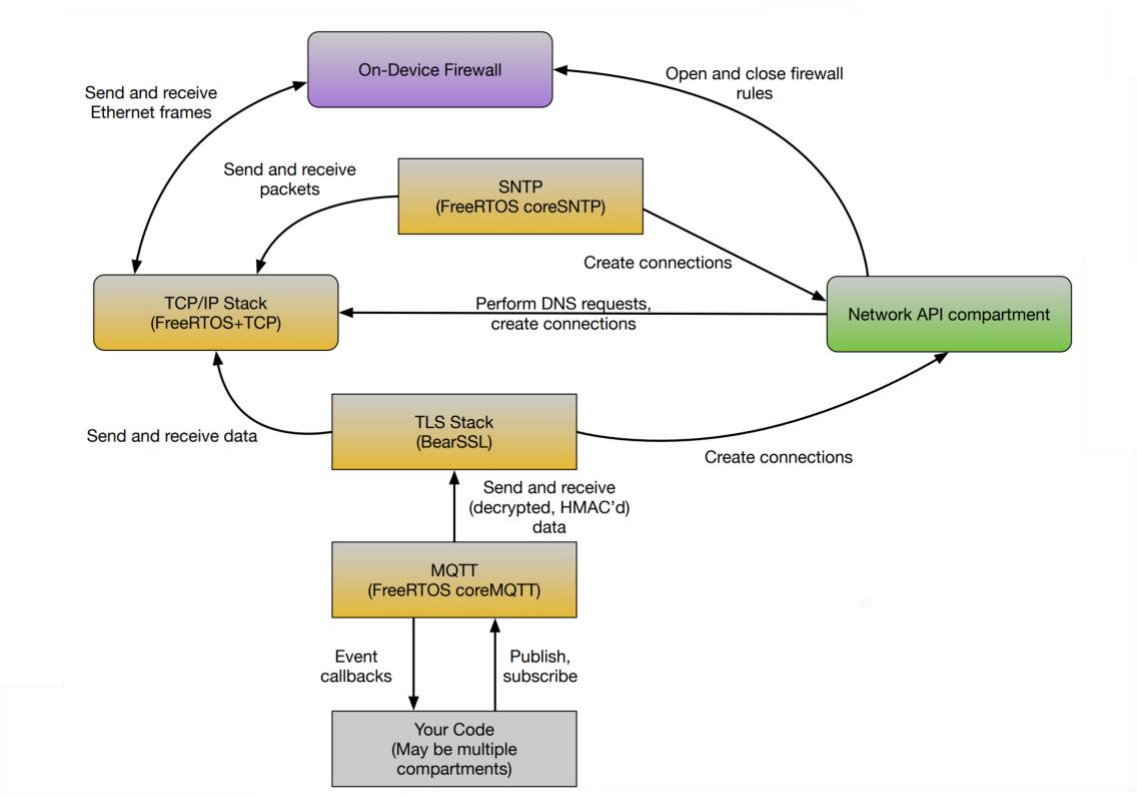

Figure 3 CHERIoT Compartmentalized Network Stack (Courtesy: SCI Semiconductor)

While a substantially simpler system than those deployed in SNbD there is a clear methodology to drive out complexity and replace it with simpler, and more testable compartments, contributing to a robust framework for understanding and evaluating constraints.

The cost of this approach was minimal, both in terms of code size impact and effort. The fundamental OSS code blocks (e.g. FreeRTOS coreMQTT and BearSSL) were untouched, and minimal effort was expended to support these as compartments. The total code size of the application grew by just 0.2% to support these new features.

Maintenance: Patching and Update Management

A major consequence of the compartmentalisation, and the ability to limit the blast radius of a CVE attack, is that organisations can now continue to operate their software with known exploits that are inherently sealed and mitigated. This broadly mitigates the consequences of zero-day attacks and enables organisations to plan and roll out updates in a structured patching and update management plan, rather than having to attempt to do rapid reactionary releases.

The cost of a zero-day patch of course changes depending on the specifics of the software, the scope of the attack, and its value. However, given the sheer amount of testing that must be done to fix a vulnerability, and the need to ensure nothing is broken by the patch, these can easily run to above $100,000 per release. The ability to impact this figure, to reduce the urgency, and ensure functionality test is spread across multiple fixes brings these costs down by an order of magnitude.

The impact on the consuming IT teams is also significant. Rather than scrambling to nail the patch into the system due to unknown consequences, the need to bring down mission-critical systems, and the burden on already over-stretched teams - the ability to resolve highly impactful CVEs into a traditional management cycle is transformational.

Emerging and Available CHERI Components

The number 1 current issue identified by global stakeholders of CHERI is the lack of commercial availability. Significant technical progress has been made around the SAIL model, FPGA, and the Arm Morello test chip, and this is to be welcomed, however, the lack of a procurable product inhibits the adoption and emboldens critics of the solution space.

CHERI, in its high processing performance form, utilising the CHERI64 architecture, has made significant progress. Most notable is the Arm Morello platform, which leverages CHERI into the baseline Arm architecture and ISA, and has enabled widespread exploration and innovation, including this SNbD project. It is to be celebrated that Arm took the courageous step to integrate advanced R&D work into their test chip and clearly understood that this was a minimal viable instantiation, which only supported core functionality. It is clear that migrating this test implementation further would require both a new v10 Arm architecture and a substantial investment of $100M's to cover the completion of the integration, verification, and software tooling impacts. As such it is understandable that Arm must wait on firm commitments from the government and significant licensees to enable further expensive progress.

In lieu of progress with Arm, the RISC-V community has made good progress in adopting CHERI into the architecture, which is currently in review by the consortium. As of today, it is likely that RISC-V based devices will be the first to market with multiple commercial

projects already underway, most notably the X730 CHERI64 IP processor from Codasip[17] and the ICENI chip family from SCI Semiconductor[18].

The Codasip X730 processor is designed to run traditional operating systems with memory-safe functionality. The design microarchitecture is 64-bit and dual-issue, enabling high clock speeds. The company's public tooling includes

- C/C++ compiler and toolchain based on LLVM17
- CHERI-RISC-V Sail model
- Das U-Boot bootloader
- Linux kernel 6.10
- FreeRTOS
- The GNU Debugger
- Yocto build system for Linux

The SCI Semiconductor ICENI device family is based on the Microsoft CHERIoT-Ibex core, a lightweight core designed for hard real-time and cyber-physical applications. The CHERIoT core is now in full version 1. release, and SCI Semiconductor has been working closely with a wide array of open source and commercial partners, including lowRISC CIC, to complete FPGA operation. This work has largely been sponsored by InnovateUK and UKRI organisations, under the guidance of the UK Department of Science, Innovation and Technology (DSIT)

The SCI Semiconductor chip solution is now progressing to a formal chip release, due in 2025, but due to the system's ability to run unencumbered on FPGA many organisations are already prototyping products utilising the platform. SCI Semiconductor and CHERIoT partners' public tooling includes

- C/C++ compiler and toolchain based on LLVM17
- CHERI-RISC-V Sail model
- QEMU open-source emulator
- Trusted Code Base bootloader
- CHERIoT-RTOS native real time OS
- FreeRTOS
- FreeRTOS Compartmentalised TCP/IP stack
- GNU Debugger
- Plus a wide variety of third-party software components.

A major difference in solutions is currently the Codasip core carries only memory-safe CHERI extensions, whereas the SCI ICENI family also integrate full compartmentalisation functionality.

[17] https://codasip.com/solutions/riscv-processor-safety-security/cheri/x730-risc-v-application-processor/
[18] https://www.scisemi.com/products/iceni-devices/

SECTION 4: Crossing the SNbD-CHERI Chasm

In product marketing, "crossing the chasm" refers to the critical challenge of transitioning from early adopters to the early majority in the technology adoption lifecycle. This concept was popularised by Geoffrey A. Moore in his book "Crossing the Chasm: Marketing and Selling High-Tech Products to Mainstream Customers."

The Technology Adoption Lifecycle

The technology adoption lifecycle is a model that describes how different groups of customers adopt new technologies over time. It consists of five segments:

1. **Innovators**: The first to adopt a new technology. They are willing to take risks and often have technical expertise.
2. **Early Adopters**: Visionaries who are quick to see the benefits of new technologies and are willing to adopt them early, despite any initial imperfections.
3. **Early Majority**: Pragmatists who are deliberate and will adopt new technologies once they see established benefits and evidence of its reliability.
4. **Late Majority**: Conservatives who are sceptical of new technologies and will adopt them only after they have become the standard.
5. **Laggards**: Sceptics who are resistant to change and adopt technologies only when absolutely necessary.

The Chasm

The "chasm" is the significant gap between the early adopters and the early majority. Crossing this chasm is crucial for the success of a high-tech product because:

- **Early Adopters vs. Early Majority**: Early adopters are willing to take risks and tolerate imperfections in exchange for the benefits of being first. In contrast, the early majority is more risk-averse and requires evidence of product reliability and value before adopting it.
- **Market Expansion**: Successfully crossing the chasm means that a product moves from a niche market of early adopters to a broader, more mainstream market. This transition is essential for achieving large-scale commercial success.

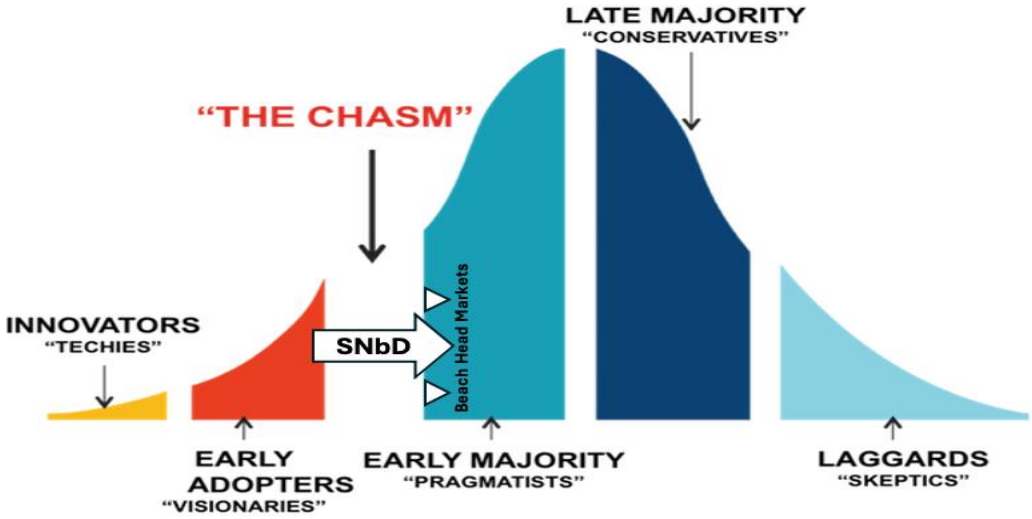

Most current implementations of CHERI, including its use in SNbD, operate within the innovation domain, supported by government-driven research and development. Commercial partners engaged with CHERI are currently engaging as "early adopters", with initial intellectual property, proof of concepts, and viable use cases emerging. To be successful with this initial transition, and the more painful transition from early adopters to the early majority, significant resources and activities must be undertaken urgently, to ensure CHERI technology has the best probability of adoption.

Recommended Strategy for Crossing the Chasm with CHERI & SNbD

1. **Targeting a Niche Market**: It is important to focus on specific niche markets within the early majority that have a clear and pressing need for the technology, enabling the construction of a strong reference base, and demonstrating implicit technical value. For CHERI there are many potential domains, but clear opportunities in areas defined in secure networking by design, immediately impacting simple routers and connected edge devices. The counterpoint to this is ensuring efforts are not spread too thinly on multiple applications, where momentum cannot be maintained for cost reasons, or where resources are sparse. The SNdB domain has already been investigated in the main work group, and it is both sensible and efficient to follow this through while paying immediate attention to the following strategies.

2. **Development of a Whole Product**: As part of the targeting process, it is obvious that the CHERI-based solutions need to complete all necessary features, support, and services that make it easy for the "early majority" to adopt and use. While investigating the SNbD domain it is clear that the solution has to primarily function as a robust router and connectivity hub, with ease of use and flexibility of implementation traditionally seen with home, small office, and simple enterprise systems. This forms a substantial barrier to entry as the SNbD team is

not an end-user vendor of these systems, and hence will need to leverage commercial partners strongly to get a minimal viable product created. Both BT and Vodaphone have been party to the group so far, and it will be interesting to see how the CHERI technology may be integrated into prototype systems.

3. **Leverage Early Adopters**: The use of success stories and endorsements from early adopters to build credibility and reduce perceived risk for the early majority is crucial in building momentum with CHERI and SNbD activities. In identifying analogous technology introductions, such as Arm winning the mobile phone processor "wars" of the early 2000s, it is important to see how the creation of push-pull marketing was required. In mobile phones Arm was successful in creating "pull" from companies such as Nokia, who then prescribed that their suppliers, the chip companies, must be based on Arm, enabling a "push" into the silicon ecosystem. For CHERI and SNdB this "pull" must be encouraged by the ultimate end-user, the UK/US governments, applying purchasing pressure on to the router and hub vendors. While this is all done openly, the router vendors need to know there is an expectation of orders (revenue) and a clearly identified new feature (CHERI) demanded.

4. **Focus on Pragmatic Solutions**: While product marketing is not explicitly an aspect of this report, it is naturally important to always highlight the practical and proven benefits of SNbD and CHERI, showing how they solve specific problems for the early majority. At one level this is simple, with the ability to reduce attacks and subsequently bring down the risk of cyber for everyone. At another, it is quite challenging as the direct purchaser of the product must understand an explicit "difference" that will impact the purchasing decision. Again, this may be regulation from the government forcing change, or it may be more implicit to the operation of the company or individual purchasing the product. In either case, the foundational impact of CHERI, with the 70%+ mitigation of critical vulnerabilities, should be sufficient to deliver value. The subsequent challenge is the viability of a "pragmatic" solution, given commercial devices of a similar performance point to modern routers are not yet available and may not be for several years. To this end, there is a need to focus on what is ready, or nearly ready, and provide solutions around this.

5. **Build a Strong Ecosystem**: "No man is an island, and no technology stands apart". We are all aware of how good technologies fail in the market due to a lack of support and ecosystem, and how the implementation of a strong ecosystem, such as mobile app stores, can turbocharge technology adoption. The same is very much true of CHERI and SNbD technologies, and while significant

resources have gone into academic research, in the next phase it is crucial to resolve a strong ecosystem of commercial partners, all of whom add their piece of differentiating technology to create the end-user solutions, often which are a long way from the core technology, and beyond the thinking space of the original inventors. This may include, for example, a distribution partner ecosystem, which is engaged on a day-to-day basis with end customers, understanding their struggles, and interpreting their future needs.

Successfully crossing the chasm involves understanding the different needs and concerns of the early majority compared to early adopters and effectively addressing them through targeted marketing, product development, and support strategies. Today, this report demonstrates that we are very much at the start of this journey, with only core technologies resolved. The remainder can evolve, with industry support over the next 3-5 years, or more.

SECTION 5: Market Analysis: Networking & Telecom

Driven by the increasing problem of cyber security, the SNbD project acknowledges the importance of the router. Routers typically sit at the "sharp-end" of the network, servicing offices, factories and homes, but due to cost sensitivity, often represent the weakest link in the chain. This is illustrated in a recent industry report that highlights routers account for over 75% of infected devices, with infected routers posing a greater threat than infected IOT devices or PCs. The SNbD project has directly addressed this threat by combining recent advances in router security (ManySecured) built on a secure CHERI computing hardware platform to demonstrate the hardened router and future networking protections.

The networking and communications industries are massive, with the telecommunications marketplace due to grow from approximately $1805.61 billion in 2023 to $3102.74 billion by 2031, at a CAGR of 6.2% from 2024 to 2031[19].
Similarly, the enterprise networking marketplace was valued at $409.3B in 2022 and is experiencing growth of over 6% through to 2027 due to the proliferation of data centres and enterprise network requirements[20].

The markets are all targets for CHERI and memory safety technologies, such as SNbD, but some are more attainable than others, primarily due to technical requirements, but also due to commercial dynamics with resistant incumbents, or the ability of governments to force change. The following tables outline the 5 biggest domains for CHERI adoption.

Enterprise Networking & Subsegments	
1. **Local Area Network (LAN):** LANs are crucial for connecting devices within a limited area, such as a building or campus. They are widely used in office environments and are essential for enabling communication and resource sharing. **CHERI Alignment** – Good. Possible MPV and alignment with CHERI64 devices.	
2. **Wide Area Network (WAN):** WANs connect devices over large geographical areas, often integrating multiple LANs. They are essential for organisations with multiple branches or remote workforces. **CHERI Alignment** – Medium. Primarily CHERI64 device targets	

[19] https://www.skyquestt.com/report/telecommunication-market

[20] https://www.globaldata.com/store/report/enterprise-networking-market-analysis/

3. **Data Centre Networking:** This involves the interconnection of data centre resources, including servers, storage systems, and networking equipment, to ensure efficient data flow and resource utilisation. CHERI Alignment – Poor due to high performance device requirements.	
4. **Software-Defined Networking (SDN) and Network Function Virtualisation (NFV):** These technologies allow for more flexible, scalable, and cost-effective network management by decoupling the control and data planes. CHERI Alignment – Medium as a secondary processor in complex systems, too early for main processor integration.	
This segment traditionally holds a significant market share, driven by ongoing investments in infrastructure, SDN, and NFV. Companies like Cisco, Juniper Networks, and Hewlett Packard Enterprise (HPE) are major players. **Estimated Market Share:** 35-40%	

Telecommunications & Subsegments	
1. **5G Networks:** The rollout of 5G technology is revolutionising mobile networks, providing faster speeds, lower latency, and the ability to connect more devices. CHERI Alignment – Poor due to high performance device requirements. Potential for Root of Trust or aux processor offload.	
2. **Fiber Optics:** Essential for high-speed internet, fibre optics support large bandwidths and are critical for modern telecom infrastructure. CHERI Alignment – Poor due to high performance device requirements.	
3. **Carrier Ethernet:** Used by service providers to offer high-speed, reliable Ethernet services over a wide area. CHERI Alignment – Poor due to high performance device requirements.	

With the rollout of 5G and advancements in fibre optics, this segment is growing rapidly. Major players include Huawei, Ericsson, Nokia, and ZTE.

Estimated Market Share: 25-30%.

Cloud Networking & Subsegments	
1. **Cloud Connectivity:** This includes services and solutions that enable businesses to connect to public, private, or hybrid clouds. CHERI Alignment – Medium/Poor due to high performance device requirements. Potential for Root of Trust or aux processor offload.	
2. **Virtual Private Cloud (VPC):** VPCs provide isolated network environments within public clouds, offering enhanced security and control. CHERI Alignment – Poor due to high performance device requirements. Potential for Root of Trust or aux processor offload.	
As cloud adoption continues to surge, cloud networking solutions are becoming increasingly important. Key players include Amazon Web Services (AWS), Microsoft Azure, and Google Cloud. **Estimated Market Share:** 15-20%.	

Industrial Networking & Subsegments	
1. **Internet of Things (IoT):** IoT networks connect sensors and devices in industrial environments, enabling real-time data collection and analysis. **CHERI Alignment –** Good. Strong alignment with simple CHERIoT frameworks.	
2. **Industrial Ethernet:** Tailored for the demanding environments of industrial operations, providing robust and reliable networking solutions. **CHERI Alignment –** Good. Strong alignment with simple CHERIoT frameworks.	

This segment is expanding due to the rise of IoT and the need for robust industrial networks. Companies like Siemens, Rockwell Automation, and Schneider Electric are prominent. **Estimated Market Share:** 10-15%	

Residential Networking & Subsegments	
1. **Home Wi-Fi Networks:** The proliferation of smart home devices has driven demand for robust and high-speed home Wi-Fi networks. **CHERI Alignment** - Good to medium**.** Potential for CHERIoT MVP or CHERI64 processor.	
2. **Broadband Internet:** Residential broadband services, including DSL, cable, and fibre, are crucial for home internet connectivity. **CHERI Alignment** - Good to medium**.** Potential for CHERIoT MVP or CHERI64 processor.	
The demand for high-speed internet and smart home devices is driving growth in this segment. Major players include TP-Link, Netgear, and Linksys. **Estimated Market Share:** 10-15%	

Industrial IoT (IIoT) Market

The Industrial IoT market is projected to grow from $194.4B is 2024 to over $286.3B by 2029[21] a CAGR of over 8.1% driven by general IoT-enabled digital transformation across all verticals, according to MarketsAndMarkets.com.
The growth in automation, enabled by the rise of AI is driving a revolution in factories. This in turn is dictating the widespread digitalisation of the workplace, with digital twins becoming standard, and communication technology becoming central to the effectiveness and efficiency of the industry.

[21] https://www.marketsandmarkets.com/Market-Reports/industrial-internet-of-things-market-129733727.html

Standardisation is also helping to drive growth, with the traditional proprietary communication architectures of operational technology giving way to IPv4 and IPv6-based communications, and both standard wired and wireless protocols in place, although extended for time-sensitive networking (TSN) or control plane applications, such as OpenCAN. Meta standards, such as those from the International Telecommunication Union (ITU) Internet of Things Global Standard Initiative (IoT-GSI), are also supporting the industry in a constant fight against fragmentation.

The IIoT marketplace covers several major vertical markets, all individually worthy of reporting. Healthcare represents the largest vertical by value at this point, however all are sizeable domains of over $10B.

Industrial IoT Markets
Manufacturing
Energy
Oil & Gas
Metals & Mining
Healthcare
Retail
Transport
Agriculture

Given our reliance on these marketplaces for our food, transport, heating, and health, these are all domains in which cyber-attacks are perilously close to bringing the world to chaos. It is commonly stated that the world is 6 days away from rioting over food shortages, and as we saw in the covid pandemic, perhaps even less for toilet rolls. As such all of these verticals are targets for memory-safe technology, with CHERI strongly applicable to separate the core function of applications robustly separated and compartmentalised from the fallible communications stacks.

ManySecured and SNbD have two roles to play in these domains.
Firstly, as a pure-play router ManySecured defined hubs enable the protection of existing systems, where devices will be connected to the network with traditional security, which is known to contain multiple CVEs. *The ManySecured hubs will form an additional layer of protection and can be used to easily replace the hubs which are currently operating, both cheaply and simply. A minimal viable ManySecured hub, primarily operating with packet forwarding, would be able to be implemented in a matter of months, once chips are available in volume.*

Secondly, once the technology is proven, these markets are dominated by simple measurement and automation endpoints, which are well suited to the deterministic performance of CHERIoT-level devices. As such this is a prime area for further investigation and investment as we look to secure our critical supply chains.

Critical Domain Vendors in this market include:

- ABB
- General Electric
- Emerson
- Intel
- Cisco
- Honeywell
- Siemens
- Huawei
- Rockwell
- PTC
- Dassault Systems
- IBM
- Robert Bosch
- NEC
- Software AG
- Texas Instruments
- KUKA AG
- Dragos
- Google
- Microsoft

Router WiFi Chipset Market

Global Market Insights state in their Wi-Fi Chipset Market – By Standard[22] (Apr 2023 GMI4849): *"There is a critical need for cyber security as a result of growing issues including wardriving, wireless sniffing, illegal computer access, theft of mobile devices, piggybacking, and other things. Regulations will be highly prioritised when enormous data sets including personally identifiable information are produced..."*

According to Future Market Insights[23] the Wi-Fi chipset market is expected to reach US$ 20.5 billion in 2023, and eventually to US$ 32.6 billion in 2033. According to the report, the market is expected to proliferate at a CAGR of 4.8% from 2023 to 2033.
Driving this growth is the continued proliferation of IoT devices, and while IoT is a superset of multiple different markets (combining consumer, industrial, medical etc.), it demonstrates that this remains a major marketplace.

The chipsets outlined here naturally flow into a wide variety of Small Office & Home Office (SOHO) applications, which is today operating primarily in an upgrade & replacement market dynamic, vs a new, quickly expanding marketplace. New Wi-Fi chipsets are being developed to be compatible with older standards to ensure seamless integration with existing network infrastructures. As chipsets continue to

[22] https://www.gminsights.com/industry-analysis/wi-fi-chipset-market

[23] https://www.futuremarketinsights.com/reports/wi-fi-chipset-market

evolve, backward compatibility will be maintained while new features will be introduced.

The connected home devices segment alone is expected to represent over 5% of the marketplace, representing in excess of $2B revenue, or roughly 500M devices p.a.

The marketplace is fragmented with several major chip vendors, including Broadcom, Infineon, MediaTek, Qualcomm and Realtek. Given a focus on impacting a broad section of the chip marketplace, it is suggested that impacting Broadcom & Qualcomm in the US, and MediaTek and Realtek in Taiwan, would deliver the highest immediate impact. Achieving this impact on their tier-1 global vendors is challenging and requires two very specific actions.

- The US and other governments must mandate progress to memory-safe solutions by 2030, or sooner. If Federal Agencies are barred from purchasing non-memory-safe technology, these large organisations will perceive a clear market advantage from supporting this activity. Some activity has been recently identified from CISA and the FBI requiring all OEMS to identify their roadmap to memory-safe implementation by 1st January 2026.
- The SNbD activities highlighted in the report must be completed and the technology proven in real-world applications. Large chip vendors are not natural risk-takers and will require substantive proof points with major OEMs.

Existing domain vendors according to Global Markets Insights (2023)

- Broadcom
- Celeno Communications
- Infineon Technologies
- Espressif Systems Shanghai Co Ltd
- GCT Semiconductor Inc.
- I&C Technology
- Intel Corporation
- MediaTek, Inc.
- Microchip Technology Inc.
- MORSE MICRO
- Newracom
- NXP Semiconductors
- ON Semiconductors
- PERASO TECHNOLOGIES INC.
- Qualcomm Technologies, Inc.
- Realtek Semiconductor Corp.
- Renesas Electronics Corporation
- Samsung Electronics Co., Ltd.
- Silicon Laboratories
- STMicroelectronics N.V.

Network Switch Chipset Market

The network switch market is projected to grow from $33B in 2023 to over $45B by 2028[24], a CAGR of over 6.5% driven by the growth of cloud computing and general data escalation, according to MarketsAndMarkets.com (Feb 2023). Chips and software form a major segment of this cost, estimated at c. $15B.

Ethernet switch chips are used to create Ethernet switches, which are devices that connect multiple Ethernet-enabled devices and allow them to communicate with each other. Ethernet switches use a variety of different technologies to manage the flow of traffic between devices, and switch chips are a key component of this technology.

There are a variety of different Ethernet switch chips on the market, each with its unique features and benefits, however, to date, Memory Safety has not been identified as a major requirement by vendors consulted for this report. The Ethernet switch chip market is largely driven by the increasing demand for high-speed networking and the continuing need for low-power consumption. The demand for bandwidth-intensive applications such as video streaming and online gaming is driving the market alongside energy efficiency.

Traditional IT security is a major requirement of the solutions, including deep packet inspection capabilities, and this will start to bring Memory Safety more into focus. Current requirements include securing remote access to inhibit DHCP port snooping and limiting MAC address learning to prevent MAC address flooding attacks.

Augmenting switch security best practices with CHERI is seen as a major win in this marketplace through the following

- Firmware updates can be managed more securely due to the inherent compartmentalisation of the codebase. It will be possible to maintain impacted code for longer without emergency zero-day patches. Furthermore, traditional memory-safe vulnerabilities are eviscerated due to inherent capabilities.
- Port security can be enhanced by placing functions into capability-bound compartments with highly restrictive software rights
- Similarly, with Access Control, the ability to strictly define capability privileges will tightly limit the ability to artificially impact these.
- Implementation of Virtual Local Area Networks through D3Controls will limit the ability of exploits to be opened up.
- Intrusion Detection Systems (IDS) & Prevention Systems (IPS) can be structured as unique compartments activated quickly out of the CHERI Trusted Code Base, ensuring defensive shields are present before the system initiates actual connectivity.

[24] https://www.marketsandmarkets.com/Market-Reports/network-switches-market-18720083.html

CHERI-based systems can deeply impact this market. However, the key trends in this space play against fundamental technology availability at this point. Increased integration is driving the need for more powerful processors to handle increased traffic loads, and hence this is driving towards traditional high-performance application processors, such as the Arm Neoverse platform.

Two solutions are available to assist CHERI adoption in this space. Firstly engagement with key vendors suggests that a multi-processor SoC solution may be viable in some use cases, where the core system can be maintained, and a CHERI processor integrated as a root of trust, or root of control, monitoring and managing larger systems. Substantial effort must be expanded to demonstrate true value differentiation and value in this system.

The second engagement is with a range of high-performance RISC-V vendors who are operating towards the performance points required for high-speed networking applications. There is a clear desire for additional differentiation, but the cost of integrating the CHERI instruction set, and verifying the implementation are very high. As such, again, it is incumbent on the major consumers of such technology to drive the need for solutions to integrate memory safety. This falls back to the government driving the network operators to adhere to its desire to progress memory safety from the top down.

Critical Domain Vendors in this market include:

- Broadcom
- Marvell Technology Group
- Intel
- Mellanox Technologies
- Arista Networks
- Cisco Systems
- D-Link
- TP-Link
- NETGEAR
- Huawei

SECTION 6: Six Obstacles to SNbD-CHERI Adoption

SNbD and ManySecured have made excellent progress to date, demonstrating the impact of Memory Safety and compartmentalisation to the networking domain. However, several clear obstacles are present that impact the ability of the industry to integrate the technology, which are listed below.

1. Industry Demand
2. Availability of Technology
3. Legacy System Integration
4. Codebase Legacy
5. Ecosystem Immaturity
6. Benchmarking
7. Market Alternatives

Industry Demand

The first and largest challenge to the introduction of radical new technology and solutions is that customers are very price-sensitive, and they must have an incredibly high motivation to make the switch to a new technology. Often the new product must hit a strong pain point, and not just be a little better, but aim for a 10x-100x improvement on critical metrics, whether performance, power, or functionality.

In security, we have a substantial industry demand challenge, in that the purchasing agent procuring IT components typically has zero connection to the cost of data breaches or ransomware attacks. While the board, or CIO, will set clear goals and objectives around security, they do not always have the ear of the finance department, at least until their brand is on fire, or they suffer a massive cyber-attack.

The mechanisms for measuring security are also challenging, as all systems are “secure” right up to the moment an exploit is found, and hence the massive issues the industry has around zero-day attacks. They are simply not priced into the purchasing decision of IT infrastructure, or technology in general.

In discussion with IT manufacturers, the solution to this challenge is threefold:

- **Vulnerability Exposure Awareness**

The first, technical, aspect of this problem is to better measure security, and enable communication around this. Typically challenging to accomplish there are evolving methods for measuring security within the Many Secured specifications including the D3 specification, which are welcome, but also the growing demands for SBoMs (Software Bill of Materials) and the enumeration of CVEs, enables organisations to now measure exposure and calculate risk.

- **Vulnerability Risk Impact**

The second is to ensure responsibility for these challenges rests with the team who own the funds to resolve it: the Board of Directors. The traditional challenge is that the CIO/CISO holds the responsibility for cybersecurity, but often lacks the resources to uphold it; this requires urgent resolution. A way to achieve this is for a full voting member of the board to assume legal responsibility for signing off corporate annual reports with specific cyber risks highlighted, and ensure cyber-insurance is valid for their organisation

- **Vulnerability Disclosure**

The third component of the solution in parallel, is for the government to ensure company regulation is tightly aligned with cyber best practices to ensure companies are mandated to disclose their true level of risk, at the same level as financial risks, both to their shareholders and their insurance providers. If this fails, formal legislation may be required.

Availability of Technology

The second most significant challenge for SNbD, as identified earlier, is the simple availability of commercial technology. Much of the project work to date has been carried out under the UKRI TAP program, using the Arm Morello test chip. This device is great for exploration but is not technically complete, and as of the time of writing, Arm has no plans to commercialise CHERI technology within their licensable IP processor portfolio. While this is certainly their right, it does create a major challenge in how the project can progress to commercial realities.

There are two alternative strategies currently that are publicly disclosed:

The first alternative is to embrace the RISC-V CHERIoT-ibex processor, which Microsoft has originated and donated to the open-source community. This lightweight processor is currently being leveraged by at least one silicon partner (SCI Semiconductor) to bring commercial products to market in the next 12 months. This is potentially a viable timescale for SNbD, especially as FPGA implementations are available today. The challenges with this approach are the performance and clock frequencies of this type of short-pipeline device are below many of the application requirements for a modern router, although viable for a minimal viable product providing packet forwarding is supported through a hardware assist engine.

The second alternative is to pause to await high-performance RISC-V 64-bit processor IP from organisations including Codasip. This technology is suitable for a wider array of performance requirements, but at time of writing is not publicly available. Furthermore, this IP is subject to licensing, design, formal verification, and substantial fabrication timelines, meaning a viable solution is a significant time away potentially 2-3 years.

The solution would therefore point to the creation of a minimum viable product (MVP) to demonstrate real-world commercial applications.

A further challenge to this availability issue is ensuring devices have either an integration baseband processor and radio, which can be expensive to integrate, or the implementation of a two-chip solution. For limited quantities, the latter is a reasonable approach, but for high volumes, a fully integrated solution would be preferable.

Legacy System Integration

As identified in the Crossing the SNdB-CHERI Chasm section, there is significant work to bring a replacement MVP router to market, with a broad range of technologies required to create a commercially viable offering. As such it is essential that the CHERI-enabled systems can build upon and leverage existing solutions, changing the parts that need to be changed, but supporting existing systems where possible, at least for an initial solution. This legacy system integration may form a problem, especially if integration is required at a SoC (System on Chip) level, where mixing CHERI and non-CHERI processors may create challenges.

Here, for example, it has been shown that CHERIoT can integrate with non-CHERI systems across the bus structure, such as TileLink, however, if shared memory is required, there is a need to mix traditional 32-bit storage, with the new 33-bit solutions required to carry the CHERI Tag bit. As we look at 64-bit CHERI there may be larger system issues, especially if the legacy systems are 32-bit.

Additional research is required to understand the scope of this issue, especially if the first commercial solution is a clean minimal design.

Codebase Legacy

Perhaps a more significant legacy issue surrounds the codebase from previous solutions that need to be protected to support SNbD applications. As we know, to our collective costs, software is widely recycled, bringing new and old vulnerabilities into our code base.

32-bit CHERIoT is a pure-capability (pure cap) programmers' model only, meaning that legacy code must be recompiled and relinked to target the new devices. This is seen as a small issue to gain access to the CHERI benefits, but it is understood that codebases get lost or polluted over time, so this is not always possible. With CHERIoT we have seen FreeRTOS network applications ported, with 0.2% of code changed to support the compartmentalisation – a fairly low barrier.

64-bit CHERI processors, when available, will support existing C / C++ code execution through a hybrid support mode, operating legacy code within a single CHERI compartment. While viable this is seen as only a first step, as most code will be improved by leveraging the CHERI compiler, which has been shown to identify Memory Safety issues, even where the code is not being transitioned to CHERI. In effect, the new compilers can be added to a modern CI/CD flow to identify hard-to-target Memory Safe flaws in C/C++, which can then be addressed directly in the code base.

Ecosystem Immaturity

As previously mentioned, ecosystems are critically important in building products and solutions. In the case of SNbD this potentially creates a near-term issue, as while significant code for CHERI exists, in fact more than RUST code today, the reality is this still represents a tiny fraction of code in the market.

Several ecosystem challenges are immediately recognised including that commercial-grade Linux does not yet exist for CHERI (vs CheriBSD). Similarly, only two Real Time Operating Systems (RTOS) have been successfully ported to the CHERIoT platform, a native CHERI RTOS, and a simple port of FreeRTOS.

Another major challenge as of today, is that the CHERIoT compiler is utilising a relatively old release of the LLVM compiler. This is a known issue, and is being addressed, but creates near-term sub-optimal performance versus the latest edition of LLVM. This is the challenge of living on the leading edge of technology, where issues will get resolved, but perhaps later than wished.

Two approaches are required to attempt to solve this issue:

1. The first is to focus in the near term on a lightweight MPV with minimum “bells and whistles”, reducing the scope for third-party code. Given the goal is to produce an inherently secure router, this minimum code support is proven to be viable.
2. The second approach is to identify the minimum viable number of ecosystem partners required to update components for the solutions needed. The impact is this level of engagement can become very expensive very quickly,

Market Alternatives

Alternative proposals to solve Memory Safety more broadly are considered here.

Firstly, as mentioned previously, Rust is often seen as a competitor to CHERI in resolving memory safety. Rust is a robust and popular language and will resolve many of the issues that CHERI does, however, it has its own challenges and requires substantial software translation, rewriting, recompiling and testing to produce a solid software platform. This effort will ensure that while new, critical, sections of code are written in Rust it is unlikely that substantial amounts of code will be ported, or github repositories resolved. It is far better to view Rust and CHERI as an “and” rather than an “or” for the real world.

Secondly, other security hardware frameworks exist, and while not directly applicable to memory safety within the processors, are supported by multiple IP vendors. Most notably IOPMP, which is a specification for a Physical Memory Protection Unit of Input/Output devices, to regulate the accesses issued from the bus masters.

IOPMP is raised to demonstrate that additional pressures are placed on industry partners who are investigating Memory Safe computing, which may derail SNbD adoption activities.

SECTION 7: The Way Forward for SNbD – A Discussion

Realising SNbD as a realistic beachhead market for CHERI technology is imperative for the project's industry partners. Firstly, as a reference design for what can be achieved with CHERI, and secondly with a secure router as a "pilot product" in its own right, impacting the networking market and protecting everything that is connected to the Internet globally. The impact of this is potentially massive, as it immediately takes out breaches, saving consumers and enterprises from phishing attacks and backing losses that impact so much of society and commerce.

However, through the research carried out for this report, the successful adoption of CHERI-based SNbD cannot be decoupled from the availability of commercially available components and tools. As a result, we have necessarily identified a set of essential requirements for next-level support. These cover government support, technology accessibility and industry partnerships. It is only by resolving these challenges that SNbD will have the opportunity to impact nationally and globally.

These essential requirements represent a "Moon Shot" for the industry, and like the Apollo mission require numerous stakeholders to act together, with robust timelines and funding, and a narrow pathway for success. This program, which we have named *Andrasta*, after the invincible god of the ancient British tribes, is challenging but represents the most likely approach to meet the next-level goals of popularizing CHERI and delivering technology globally.

Andrasta SNbD Moon Shot Program

Ensure Continued Government Leadership & Support ○ 5-year commitment ○ £80M est.
Continued governments' support is not guaranteed and must be shown as value for money against other public priorities. However, as a bedrock capability to all governmental functions and networking the protection of IT and OT systems are inherently central to driving efficiency. Inherent memory safe network operation is valuable, but if, through Many Secured, we can once again enhance the lifetime of existing IT/OT system then we potentially reduce the cost of rebuilding government infrastructure by many billions, making this investment tiny given the scale of national IT spend.

Four specific outcomes are desirable.

a) Deliver Public funding & support for goals
The first requirement from government is to create broad alignment for ManySecured as a natural go-to-market solution for the CHERI technology both as a technology demonstrator for MemSafe technology, and as an end in its own right.

This commitment to the project needs to actively encourage key equipment vendors to the government to embrace the technology and support the further development of the project. Notably BT & Vodaphone are already engaged with ManySecured, but transforming this from an early technology demonstrator to something that is deployed to people's homes is complex and multi-faceted.

It is further critical that the UK government work with partner nations US/Aus/Can/NZ, Singapore and beyond, to support the evolution of the technology, to ensure we do not become technologically isolated, and support market uptake.

b) Impact structured purchasing requirements for Memory Safe technology.
Traditionally the UK government, and its various organs, set supplier requirements, but stopped short of defining how functions must be delivered. For example, one active security group within the government has stated they want to see more systems "Memory Safe", but has publicly stated they do not wish to decide how this is achieved.

While the sentiment for this is understandable, the reality is that this leaves the industry floundering, investigating multiple dead-end solutions that academia closed off many years ago.

In the US the government can set very tight definitions of technology to be supplied to Federal Agencies, driving very narrow and explicit behaviours, and impacting markets overnight. *It is strongly recommended that the UK government follow this path, instructing critical stakeholders, such as NCSC, Ofcom and Ofgem, to mandate hardware enforced Memory Safe technology within 5 years in all communication systems.*

c) Ensure cyber posture is explicitly required in corporate annual reports
From an SNbD, and CHERI, perspective, companies must be persuaded to move to a cyber-first approach, in the same way that safety-first has been for the past three decades.

To impact corporate culture the government has supported board level inclusion targets for gender, and more modern requirements for ESG (environmental, social and governance*). It is strongly suggested that cyber is added to this list, to ensure annual reporting, and day-to-day activities around cyber resilience are integrated into normal business flow.* Specifically, this work needs to ensure cyber moves out of explicitly referencing IT (or OT) systems, and instead becomes part of the fabric of

business, impacting product and service offerings, supply chain discussions, financial planning, and employee training.

d) Deliver extended support for supply chain

Delivering commercial grade ManySecured and SNbD infrastructure is a laudable goal, but as the saying goes "it requires a village to raise a child". SNbD will only be successful if the supply chain into the projects is viable and supported, starting with the technology vendors. Existing government spending on CHERI has been substantial and very welcome, however as we progress from pure R&D to commercialisation some UK plc 'leaders' need to be supported through the next stage of technology adoption. Who chooses the leaders is outside of the remit of this report, however it is clear that a whole new technology industry must be supported to enable Memory Safe technology. Current solutions support the status-quo too heavily, and gaining traction requires demonstrable progress, driving a catch-twenty-two cycle of "solutions are not ready, so we can't start moving".

Foundational supply chain components which require additional investment and support include IP generation, chip development, tools development (compiler, debugger, OS, etc), with the ability to deliver real world applications.

Deliver Near Term Solutions Based on Available Technology

- 3-year commitment
- £10M est.

While development on the Morello test chip has been very positive for SNbD activities, the unfortunate reality is that it is not commercially available or relevant. *As such a decision must be made on how to proceed, and it is suggested that SNbD look at developing a minimum viable product based on the CHERIoT core, which is mature and available, and is gaining moment as it moves towards a commercial reality.*

The core is relatively simple, versus the Morello Neoverse test chip, but has advantages because of that, with achievable implementation and faster time to market. Being a simple device, it is suggested that SNbD look at how this chip can be utilised in a larger system, alongside existing chip sets, to provide a secure central core inside a traditional networking framework. While this solution is not perfect and will not demonstrate all of the intended functionality of SNbD, it will enable CHERI to "augment" and extend existing solutions, bring down barriers to entry, and enabling a good, better, best set of solutions to emerge. In this way the transition to CHERI enabled systems is also broken down, reducing complexity and risk for existing communication device vendors, creating a viable beach head, and demonstrating technology in a far shorter period.

For this report substantial initial investigation has taken place into aspects of porting SNbD onto a CHERIoT platform. Providing focus is primarily focused on packet forwarding, with secure D3 management operating on the CHERIoT processor, it is viable to build a high bandwidth (but sub-GHz) system, where table lookup, port destination and header updates can be built, firewalling the system and providing a lightweight router solution.

Enable Mid Term Applications Based on Evolving Technology

- 5-6 year commitment
- £20M-30M est.

The CHERIoT processor exists in FPGA, and is moving towards silicon quality solutions, but is limited by the performance a short central processing pipeline can achieve. To fill the performance gap between the CHERIoT core, and the performance seen with the Morello test chip it is important to impact the design and implementation landscape now, to deliver the desired solutions of tomorrow.

It is suggested that once the CHERIoT device is in fabrication, that funding be sought to deliver a more complex communication SoC based on CHERI64 solutions, with activities similar to those being currently completed by lowRISC, on behalf of UKRI. The complexity of this SoC is at least an order of magnitude greater than that of Sunburst, and as such at least £20M of funding would be required to take this forward.

This is a substantial investment, and something government funding needs to look at carefully. However, an alternative is available through a specialist public-private partnership, where if the government commits support to the project, with a level of investment, then VCs and other traditional investment can be sourced to complete the fundraising. In this manner the government, or one of its agencies, becomes a major shareholder, seeing potentially significant returns on investment and support, plus receives the outcomes it desires. There are governmental engaged VCs, such as NSSIF, which are potentially good vehicles for this, however, to drive this program explicit government leadership from DSIT, Cabinet Office, MoD, and Treasury must be identified and attach resources.

Move Fast and Break Things

- 3-5 year commitment
- £10M est.

The ManySecured program has already provided valuable feedback and should be seen as a strong success, limited only by factors outside its control (Morello test chip availability),

As mentioned earlier available technology is limited in the performance space required for routers and hubs, and while a hybrid solution of a CHERI central core

controlling a more traditional communications system is not ideal, it is certainly viable.

An alternative to this solution is to demonstrate how Memory Safe technology can be applied to the simpler edge nodes of the system, for example impacting specific IoT measurement and control points that are themselves connected back to the routers. This is fully in line with expected behaviour and functionality of the CHERIoT based devices, and builds memory safe applications "from the outside-in". The nature of these IoT edge devices is they are also functionally simpler to construct, faster to gain entry to market, and subject to far less purchasing scrutiny. In many aspects these offer a "free hit", as if they do not meet expectations they can simply be exchanged.

It is suggested that the SNbD project identify a critical communication-centric application, for example the development of a simple transportation monitoring and telematics system, to investigate how the critical learnings from the project should be applied to devices servicing this market. If it is possible to provide security at the edge node, similar to what is envisioned for the gateway then achievable market impact is potentially very near.

Move Fast and Fix Things

- 1-3 year commitment
- £5M est.

While it may not be possible to resolve an entire router with a simple CHERIoT processor, it may be viable to create a set of data-diode "dongles" which provide much of the router security values of identification, isolation and device management in a distributed form. Sitting between the router and the IoT device as a wifi "hop" or on a wired interface, it is possible to envision these devices being cheaply and easily rolled out as power socket blocks which connect to the main wifi router for their connectivity, but then service the explicit device on the far side, for example a connected refrigerator, printer, or connected lighting. This solution may be limited on connectivity throughput, but would act as an isolation layer in the home. Traditional routers, such as the BT HomeHub, could be extended to recognise and support these dongles, providing a simple in-home experience for the consumer.

Conclusions

ManySecured and SNbD have been highly successful projects to date, highlighting the need for next-generation networking schemas built on the CHERI platform as a demonstrator for the memory-safe revolution. The market and commercial challenges around technology maturity and availability weigh strongly in the report and there remains a significant effort to ensure this frontier technology does not get lost in the technology chasm.

Ultimately a call to action is required to ensure the future ambition of fixing our digital foundations is achieved. This will be accompanied by sizable and sustainable social and economic benefits. There is only so much a single project can achieve on this journey, yet we have helped demonstrate CHERI's potential. As demonstrated in this report, the industry is on the cusp of driving memory-safe technology forward yet requires follow-through investment and procurement support from government to complete the job of ensuring the industry pivots over the next 5 years. Now, we must further leverage the significant work of UKRI and the Digital Security by Design Programme, with forward-looking views of the UK government and major US stakeholders including DARPA, CISA and the FBI.

To achieve the needs of our increasingly digital and connected world, critical recommendations beyond the project include:

- Targeted investment to drive device availability.
- Ensuring continued government leadership & support.
- Deliver near-term solutions based on available technology.
- Continued investment into academia and ecosystem development.
- Supporting the transition to memory-safe devices based on CHERI/CHERIoT.
- Implementing the UK equivalent to Federal Purchasing requirements.
- Exploring the role of Cyber Insurance to include Memory Safety incentives to help create market pull.

Annex I - CHERI Picking Applications

CHERI is a welcome and game-changing technology, yet it faces many new technology adoption challenges; the current lack of commercial device availability, finite device performance (in the near term), yet with the desire to impact an entire industry. It is therefore important to apply the principles of New Product Introduction (NPI) and identify what can be done to win beachhead markets and identify the resources to follow through beyond the beachheads to broad market adoption.

Driven by Need - Industrial Segment Review

The first action is to identify a clear set of industrial sectors where network application security, integrity, and, specifically, memory safety are not merely justified, but strongly required. This work has evolved through discussions with a wide variety of stakeholders across government, industry, and academia:

Sector	Sub-segment identification (critical markets in **bold**)
Manufacturing & Automotive	**Automotive**, electronics, machinery, textiles, consumer goods, etc.
Energy	**Electricity** (**generation**), oil and gas, **renewable energy** (solar, wind, hydro), nuclear power, coal, etc.
Healthcare and Pharmaceuticals	Hospitals, **medical devices**, pharmaceuticals, biotechnology, healthcare services, etc.
Telecommunications	Fixed-line, mobile, satellite, internet service providers, **telecom equipment,** etc.
Information Technology (IT)	Software development, **hardware manufacturing**, IT services, data centers, **cybersecurity**, etc.
Utilities	**Electricity**, **water**, natural gas, waste management, etc.
Aerospace and Defense	Military aircraft, commercial aircraft, spacecraft, **defense equipment**, etc.

Figure 4 Major aligned industrial sectors

While all segments and industries remain price-conscious, it has been noted that the segments and domains identified in Figure 4 are more willing to balance the cost and benefit of memory-secured devices, versus, for example, retail and consumer goods. This is critical as new products will likely have a price premium to cover the cost of upskilling and development. Another notable aspect is that most of these segments are also regulated through government or independent industry bodies, ensuring the public is protected and that the entire sector's behaviour meets expected levels of common good.

In the Energy and Utility domains, for example, the need to provide cheap, reliable energy to the public and industry, means clear guidance on system behaviour, ongoing investment and network uptime. Similarly, the fundamental push towards energy efficiency, encompassing cleaner generation and distribution of energy, plus sizable electric vehicle (EV) charging demand, is driving strong discontinuities which require

substantial investment. This in turn is enabling the evolution of newer and smarter systems, including CHERI-based devices.

Networking & Communications

As highlighted in this report, SNbD - Secure Networking by Design - is critical in all of the identified sectors.
The balance made by organisations' in each of these sectors is between ***cost of system purchases****,* ***lifecycle management****,* ***compatibility*** *with existing infrastructure, and the underlying* ***functional behaviour.*** *Upcoming threats and the need to protect* ***critical resources****, manage* ***control & operational points,*** *secure* ***intellectual property****, and/or safeguard* ***brand*** *and* ***reputational damage*** *have been identified as critical, once purchase and integration cost requirements have been met.* It would be welcome for this decision tree to be inverted, where the potential impact of an attack drives changes in purchasing behaviours - but our research indicates we remain substantially away from that situation for now, and at least until specific legislation or cyber insurance regulation is introduced.

CHERI-Solutions Market Adoption: Phases & Targeting

This annex details a temporary – yet addressable - mismatch between industry capabilities, technology performance, device availability and customer needs for CHERI-enabled systems. This 'chasm' is a natural part of the technology's evolution and should be viewed as an opportunity for early adoption rather than a long-term obstacle. There is undoubtedly a significant challenge to achieving CHERI-enabled implementation of high-end server and cloud devices, but the most likely way to address the investment need is to demonstrate security capabilities on the systems we have in the near term, and ultimately demonstrate to high-performance device IP vendors, and integrated system-on-chip (SoC) developers that it is worth investing the sizable investments required to cross this chasm.

The solution identified by this research is evolutionary and requires more than a product development plan: an industry needs to be created to build integrated CHERI implementations based on the technology available today, to demonstrate success in specific beachhead markets, and to build a viable ecosystem of code and components that can act as the bedrock for future device evolutions. If this can be achieved, it will fundamentally fix the foundations of computing to ensure future connected systems are secure and fit-for-purpose.

The following plan sets out how this could be achieved with the SNbD project as an example in three phases: near-term 'enablement', mid-term 'evolution' and long-term 'differentiation'.

Near-Term Enablement

The performance requirements for networking are highly dependent on the application and the tasks required. Some networking applications require extreme processing of multiple cores, threads, and GHz, with very high system throughput, speculative execution, multi-channel management, and so on. Today's CHERI-enabled processing cores and devices are significantly lower-power and simpler than these heavy-lifting high-performance network devices but do offer significant advantages in lower-end applications.

The Microsoft CHERIoT-ibex[25] processing core is a high-efficiency 3-stage pipeline processor, developed for robust deterministic applications, such as cyber-physical control systems, industrial and consumer applications.

In the networking and communications domain, this device has already been integrated into systems with the support of the UKRI (UK Research & Innovation) Sunburst program, in collaboration with various organisations including lowRISC and SCI Semiconductor. Predominantly delivered as an FPGA implementation, SCI Semiconductor is collaborating across the industry to deliver a physical device in 2025 which will support entry-level communications including Ethernet with TCP/IP protocol. While specific benchmarks have not yet been publicly released the device has shown to support robust CHERI-based memory safety and compartmentalisation of the FreeRTOS network stack, resolving all known CVEs associated with this software. Figure 5 below demonstrates ten known critical vulnerabilities with the FreeRTOS network stack resolved through the CHERIoT processor instruction set architecture in combination with compartmentalisation.

[25] https://github.com/microsoft/cheriot-ibex

CVE	Description	CHERIoT ISA	CHERIoT Compartment
CVE-2018-16522	Amazon Web Services (AWS) FreeRTOS through 1.3.1 has an uninitialized pointer free in SOCKETS_SetSockOpt.	■	
CVE-2018-16526	Amazon Web Services (AWS) FreeRTOS through 1.3.1, FreeRTOS up to V10.0.1 (with FreeRTOS+TCP), and WITTENSTEIN WHIS Connect middleware TCP/IP component allow remote attackers to leak information or execute arbitrary code because of a Buffer Overflow during generation of a protocol checksum in usGenerateProtocolChecksum and prvProcessIPPacket.	■	
CVE-2018-16525	Amazon Web Services (AWS) FreeRTOS through 1.3.1, FreeRTOS up to V10.0.1 (with FreeRTOS+TCP), and WITTENSTEIN WHIS Connect middleware TCP/IP component allow remote attackers to execute arbitrary code or leak information because of a Buffer Overflow during parsing of DNS\LLMNR packets in prvParseDNSReply.	■	
CVE-2018-16599	An issue was discovered in Amazon Web Services (AWS) FreeRTOS through 1.3.1, FreeRTOS up to V10.0.1 (with FreeRTOS+TCP), and WITTENSTEIN WHIS Connect middleware TCP/IP component. Out of bounds memory access during parsing of NBNS packets in prvTreatNBNS can be used for information disclosure.	■	
CVE-2018-16524	Amazon Web Services (AWS) FreeRTOS through 1.3.1, FreeRTOS up to V10.0.1 (with FreeRTOS+TCP), and WITTENSTEIN WHIS Connect middleware TCP/IP component allow information disclosure during parsing of TCP options in prvCheckOptions.	■	
CVE-2018-16527	Amazon Web Services (AWS) FreeRTOS through 1.3.1, FreeRTOS up to V10.0.1 (with FreeRTOS+TCP), and WITTENSTEIN WHIS Connect middleware TCP/IP component allow information disclosure during parsing of ICMP packets in prvProcessICMPPacket.	■	
CVE-2018-16600	An issue was discovered in Amazon Web Services (AWS) FreeRTOS through 1.3.1, FreeRTOS up to V10.0.1 (with FreeRTOS+TCP), and WITTENSTEIN WHIS Connect middleware TCP/IP component. Out of bounds memory access during parsing of ARP packets in eARPProcessPacket can be used for information disclosure.	■	
CVE-2018-16602	An issue was discovered in Amazon Web Services (AWS) FreeRTOS through 1.3.1, FreeRTOS up to V10.0.1 (with FreeRTOS+TCP), and WITTENSTEIN WHIS Connect middleware TCP/IP component. Out of bounds memory access during parsing of DHCP responses in prvProcessDHCPReplies can be used for information disclosure.	■	
CVE-2018-16603	An issue was discovered in Amazon Web Services (AWS) FreeRTOS through 1.3.1, FreeRTOS up to V10.0.1 (with FreeRTOS+TCP), and WITTENSTEIN WHIS Connect middleware TCP/IP component. Out of bounds access to TCP source and destination port fields in xProcessReceivedTCPPacket can leak data back to an attacker.	■	
CVE-2018-16523	Amazon Web Services (AWS) FreeRTOS through 1.3.1, FreeRTOS up to V10.0.1 (with FreeRTOS+TCP), and WITTENSTEIN WHIS Connect middleware TCP/IP component allow division by zero in prvCheckOptions.		■
CVE-2018-16598	An issue was discovered in Amazon Web Services (AWS) FreeRTOS through 1.3.1, FreeRTOS up to V10.0.1 (with FreeRTOS+TCP), and WITTENSTEIN WHIS Connect middleware TCP/IP component. In xProcessReceivedUDPPacket and prvParseDNSReply, any received DNS response is accepted, without confirming it matches a sent DNS request.		■

Figure 5 CVE Vulnerabilities Resolved in FreeRTOS Network Stack with CHERIoT Processing Core + CHERIoT RTOS

CHERIoT-based devices are able to support a wide variety of simple networking requirements, as defined by SNbD, as will be explicitly identified below.
For clarity, it is believed that this core, and near-term devices, will enable connectivity across a range of industrial segments including:

Manufacturing & Automotive	- Simple (MVP) CHERI enabled automotive **ManySecured** router / gateway - System comms data-diode dataflow management - System root of trust (RoT) functionality - Simple automotive telemetric connectivity - Secure general-purpose communications processor alongside baseband processor - Execution of simple communication protocols (Matter, etc) - Execution of operational technology protocols including industrial ethernet, fieldbus and wireless - Execution of SCADA protocols including DNP3 & IEN 60870-5-104
Energy	- Simple (MVP) CHERI enabled industrial **ManySecured** gateway - System comms data-diode dataflow management - Data-diode for legacy OT systems - Low voltage control & communications - Smart metering to distributer communications - In-home smart energy white goods - EV and power-wall communication & control - Dynamic pricing control & communications

	- Renewable-energy sourcing control & communications
Healthcare and Pharmaceuticals	- Simple (MVP) CHERI enabled industrial **ManySecured** gateway - Data-diode device for legacy IT systems (hospital, remote health, etc) - System root of trust for legacy IT systems - Communication for home support services - Data management for cross-domain systems - Medical device control and communicate (e.g. insulin pumps, alarms, etc) - Local networking & routers in-premise
Telecommunications	- Small office, home office (SoHo) router based on simple CHERI enabled **ManySecured** gateway - Data-diode devices for legacy telecoms systems (additional security dongles) - Wireless networking hop points - After-market root of trust services - Information flow and domain filtering - Exfiltration monitoring and blocking - Simple firewall applications - DNS services - Secure network access points (w. baseband device)
Information Technology (IT)	- Simple (MVP) CHERI enabled commercial **ManySecured** gateway - Secure WiFi points (w. 802.11 device) - Data-diode devices for legacy IT systems - Secure network attached storage - Secure network printing
Utilities	- Simple (MVP) CHERI enabled industrial **ManySecured** gateway - Data-diode devices for legacy OT systems - Network access points for OT systems - Secure router for OT systems - Secure networking of OT (w. baseband radio) - Network connectivity for metering - Network connectivity for distribution & mgmt.
Aerospace and Defense	- Simple (MVP) CHERI enabled mil-enabled **ManySecured** gateway

	- Data-diode for low-bandwidth x-domain mgmt. - Data-diode for legacy mil-aero systems - Communication networking for drone mgmt. - Modular secure networking in theater systems - Secure communication and control systems

Two additional major limitations are insight at the point of report writing.

1. **RTOS:** Firstly, from an operating systems perspective, the CHERIoT device only currently has Real-Time Operating System (RTOS) support, with native CHERIoT RTOS and FreeRTOS operating on the processor. This is seen as a natural starting point, however the implementation of a richer operating system, such as Embedded Linux would be a welcome addition. CHERI eradicats the differentiation between Memory Protection Unit (MPU) and the more performant Memory Management Unit (MMU), and hence it is possible to operate lightweight Linux type operating systems even on smaller processors.

2. **Device Integration:** Secondly, from the device perspective, while these are easy, simple, devices which support common ethernet communication protocols via a simple SPI control interface, there is a potential lack of integrated baseband radio communications. Again, as market demand grows this is likely to be resolved, but in the meantime, simple low-cost dual-chip solutions are viable to resolve the market and deliver CHERI technology into key beachheads.

Mid-Term Evolution

The inability to progress the Arm Morello test chip to a commercial reality is a challenge for the organisations which have leveraged it widely through the UKRI Digital Security by Design (DSbD) Technology Access Program (TAP). Over 40 organisations have progressed through the TAP, with significant opportunities identified and novel technology exploitation.

While the Morello test chip is based on an Arm Neoverse architecture device, which is focused on enterprise and infrastructure workloads, the architectural roots go back to the more traditional Cortex-A processors, which are optimized today for mobile and high-end embedded application. These Cortex-A processors are broadly equivalent to the majority of RISC-V implementations, and hence it is forecast that in the mid-term we will see several RISC-V mobile and embedded processors adapted to utilize CHERI technology.

The most advanced of the RISC-V CHERI application processor implementations today is that from Codasip, with their Codasip-700 family of processors, although most details remain confidential at the time of authoring. More information is available on the Codasip website[26].

The nature of IP licensors is that they do not themselves create the end devices, and licensors must acquire the IP, integrate a system-on-chip around it, and ultimately have it manufactured. Hence even when the IP is released there is a significant journey – typically measured in years - before physical silicon availability. Beyond this, there is a clear need for copious amounts of software, not least a commercial grade operating system such as Linux. Variants of open-source operating systems are available and have been developed by a collaboration of universities led by the University of Cambridge, with CheriBSD being the most notable. Significant effort has been invested by a wide collaboration of universities and commercial partners around the Morello test chip to date and while a commercial solution is not “oven ready” right now, once commercial hardware is released, we expect to see a broadening ecosystem within two to five years.

From an SNbD perspective, it may be anticipated that sufficient pickup of many of the core requirements based on robust, 1GHZ+ CHERI-enabled processors, will enable the ability to build a fully CHERI-enabled router, or ManySecured, device. It is further expected that while application processors will exist independently, there will be a rapid push toward processors with integrated communications.

The table below illustrates the mid-term potential for each segment.

[26] https://codasip.com/solutions/riscv-processor-safety-security/cheri/

Manufacturing & Automotive	- CHERI enabled automotive **ManySecured** gateway - Automotive telemetric connectivity - V2X (vehicle-2-X) communications - Secure general-purpose communications processor - Implementation of operational technology protocols - Advanced secure networking points
Energy	- CHERI enabled industrial **ManySecured** gateway - Fully integrated smart metering systems - Integrated LV/HV control & communications - Dynamic access and pricing of energy markets
Healthcare and Pharmaceuticals	- CHERI enabled medical-grade **ManySecured** gateway - Data management for cross-domain systems - Replacement of legacy IT systems (hospital, remote health, etc.) - Communication & control of home support services - Advanced Medical device control and communication - Replacement secure networking
Telecommunications	- CHERI enabled **ManySecured** gateway - Mid-range networking & router appliances - Wireless networking access points - Next-gen firewalls & domain filtering - Network monitoring and control -
Information Technology (IT)	- CHERI enabled automotive **ManySecured** gateway - Secure WiFi points (w. 802.11 device) - Advanced network attached storage - Advanced network printing - Secure signage - Systems monitoring and management
Utilities	- CHERI enabled industrial **ManySecured** gateway - Advanced network access points for OT systems - Advanced secure router for OT systems - Integrated secure networking of OT - Advanced connectivity for metering - Advanced connectivity for distribution & mgmt.

Aerospace and Defense	- CHERI enabled military grade **ManySecured** gateway - Secure high-integrity communication networking - Integrated secure networking in theater systems - Secure communication and control systems - X-domain (secure-to-internet) communications

Major challenges to the mid-term roadmap reflect the reality of "crossing the chasm" and the need for the industry to gain traction before further/major investment is likely to be released into the industry.

This leap forward relies on successfully driving a small number of beachhead markets, potentially with near-term solutions, based on CHERIoT, or waiting a (hopefully small) number of years for the mid-term technology to arrive with a mature ecosystem.

Several challenges are specifically identified for the mid-term:
Mid-Term Challenge 1: In the search for a "better" application processor solution, the industry will likely need to embrace the "good" or "good enough" near-term solutions. This is to ensure that the technology is continuously proven and organisations are seen to be successful with beachhead application integration.

Mid-Term Challenge 2: In parallel, the industry must solidify the ecosystem, and ensure it is not "chasing the next shiny object" – that is, it must remain focused and objective. Academia and open source are fantastic resources, yet for commercial success we need applications and code libraries to be completed be able to get to market, and critically, to be maintained. This is especially true of compilers, debuggers and operating systems, which form the bedrock of the ecosystem. For example, formal compiler support must be implemented in LLVM18 both to gain from substantial optimisation and to ensure developers can implement code swiftly.

Mid-Term Challenge 3: Given mid-term requirements for rich and robust Linux operating systems, there is significant heavy lifting to be carried out across the ecosystem. It is estimated that Arm and their partners invested over $100M into the Linaro organisation to get to an optimized Linux variant for the Arm architecture. Whether the CHERI ecosystem has equivalent resources is outside of this report, but it does require large enterprises with sufficient budgets to step up to consume the technology and drive an initially imperfect solution (see Mid-Term Challenge 5).

Mid-Term Challenge 4: While it is known that many large global enterprises are investing, investigating, and supporting CHERI, such as Microsoft with the CHERIoT processor, there is a clear need for the "cloud giants" and other critical organisations to voice support for CHERI to ensure mid-term success. In a classic catch-22, unless large organisations demand memory-safe hardware, there will be less appetite from investors to support the young companies currently driving innovation.

Mid-Term Challenge 5: Finally, and in support of the fourth requirement, governments in the UK, US and globally, who understand the importance of memory safety, should start to mandate it in their requirements specifications for next-generation systems. For example, in the UK it is suggested that next-generation smart meters mandate memory safe technology be used for communications and control. And in the US, it is suggested that the Presidential Order and Cybersecurity Improvement Act be updated to require that Federal Purchases mandate memory-safe technology, especially around secure networking. These two simple acts would encourage large organisations to address the memory-safe challenge, creating demand and encouraging the development of the CHERI ecosystem. The authors acknowledge and appreciate recent supportive statements from the CTO of NCSC in their CyberUK'24 keynote, and the broad support of the Office of National Cyber Director (ONCD) at the White House for their 2024 memory-safe report.

Long-Term Differentiation

It is important to have a clear vision for CHERI technology, and a roadmap to achieve it. From the author's perspective, there is no reason why CHERI cannot be adopted across the entire span of processors, from simple microcontrollers and single-board computers, all the way to the most complex enterprise and cloud systems; and from modern RISC-V processors, through Arm, and into the traditional x86 domain. Modern CISC (complex instruction set computers) are - *in reality* - RISC cores with microcode, so whilst challenging, the opportunity to impact these is possible.

In reality, we are almost there - the Arm Morello test chip is based on a Neoverse platform, and while this may not be the most modern implementation and the entire architecture has not been modified to support CHERI, it has demonstrated that CHERI can be implemented in a high-complexity system and points to a bright future.

As always, the challenge is cost. It is estimated that the cost to develop CHERI-enabled IP and validate it on a high-performance RISC-V architecture, for networking-level applications, would be at least $50M and likely closer to $100M. While this is a notable amount, it is well inside the ability of large organisations to fund, or for venture capitalists, and/or governments, to support. However, significant proof points are required before any organisation would be willing to take this level of risk. CHERI must be proven in the market, the ecosystem must be demonstrable, and the customers must be requesting/demanding the technology for real-world applications.

The long-term differentiation of CHERI-enabled applications has the following potential:

Manufacturing & Automotive	- Intrinsically memory safe automotive platforms - Intrinsically memory safe self-drive + ADAS - Memory safe robotics - Memory safe platooning of vehicles - Memory safe integrated manufacturing flows
Energy	- Memory safe power generation - Memory safe power distribution management - Dynamic pricing to consumers based on system loading and renewable generation - Delivery of efficient, next-generation power grid
Healthcare and Pharmaceuticals	- Removal of ransomware risks in healthcare - Robust data management for cross-domain systems - Next generation IT systems for integrated care (hospital, remote health, etc.) - Advanced medical robot control and communication

Telecommunications	- High performance networking & router appliances - Secured wireless networking with mass virtual network overlay - Advanced firewalls & domain filtering - Advanced network monitoring and control
Information Technology (IT)	- Secure networking - Advanced system monitoring and management
Utilities	- Advanced networking for OT systems - Advanced connectivity for metering systems - Advanced control of distribution & mgmt.
Aerospace and Defense	- Secure high-integrity communication networking - Integrated secure networking in theater systems - Secure communication and control systems - X-domain (secure-to-internet) communications

As with the mid-term requirements, there is much to be celebrated, with significant ecosystem formation and industry liaisons having been forged, however given the costs and a need to focus on mid-term deliverables this work needs to evolve over the next 2-5 years.
To achieve traction in this marketplace a number of aspects need to align.

Firstly, CHERI must be proven in the market in the short and medium term. Even though the performance requirements are widely differentiated, all stakeholders are looking for market data around performance, security impact, and market acceptance.

Secondly, the foundations of the ecosystem must be further developed, to minimize risks around adoption, and ensure the solution space is correctly bounded. No company wishes to invest $100M and then discover they need to further create the ecosystem.

Third, in this context, no single organisation has expressed a desire to be the first mover. While this could be an advantage, if only one vendor offers CHERI-based solution, then end customers may not consume the product for risk of lock in. As such there is a calculation to be made as to how and when to be the first mover within a raft of secure solutions.

Forth, there is a tremendous amount of legacy technology in the system and the transition needs to be carefully managed – from first movers to laggards. The CHERI-enabled hardware must seamlessly integrate into legacy IT systems and immediately enhance the security posture. Furthermore, the legacy code base needs to evolve. Some of the kernel components may be ported to RUST, but much of the legacy code

will need to be recompiled into compartments. This should be relatively low risk however work on this amount of code is not simple, or cheap.

Fifth, the end users, who are predominantly cloud giants and large international telecoms vendors, need to be certain that the solutions they look to integrate will be accepted by the ultimate stakeholders, the governments of the countries they operate in. At present, the US and UK governments, alongside Australia, Canada and New Zealand, have publicly called for memory-safe infrastructure build-out. At some point, the governments will need to back this up with infrastructure contracts that mandate the technology is present.