# Code Migration & Memory Safety

CRESP Industry Day, November 14, 2024 Per Larsen, Immunant, Inc.

#### Who?

- Immunant, Inc.
  - Code migration from C to Rust
  - Prevention of memory corruption vulnerabilities
  - Rust training 🦀
- Per
  - CEO at Immunant, Inc.
  - Postdoc at University of California, Irvine
  - PhD from Technical University of Denmark 🚼





https://www.cisa.gov/resources-tools/resources/product-security-bad-practices

# **Development in Memory Unsafe Languages (CWE<sup>[1]</sup>-119 and related**

Development in Memory Unsafe Languages (<u>CWE<sup>11</sup>-119</u> and related weaknesses)

The development of new product lines for use in service of critical infrastructure or NCFs in a memory-unsafe language (e.g., C or C++) where there are readily available alternative memory-safe languages that could be used <u>is</u> dangerous and significantly elevates risk to national security, national economic security, and national public health and safety.

## BADMER

For existing products that are written in memory-unsafe languages, not having a published memory safety roadmap by January 1, 2026 is dangerous and significantly elevates risk to national security, national economic security, and national public health and safety. The memory safety roadmap should outline the manufacturer's prioritized approach to eliminating memory safety vulnerabilities in priority code components (e.g., network-facing code or code that handles sensitive functions like cryptographic operations). Manufacturers should demonstrate that the memory safety roadmap will lead to a significant, prioritized reduction of memory safety vulnerabilities in the manufacturer's products and demonstrate they are making a reasonable effort to follow the memory safety roadmap.

#### Memory safety

- Means that software or a programming language is designed to catch memory bugs

#### Memory safety helps avoid

- Security Vulnerabilities
  - 60-90 percent of software flaws in memory unsafe languages
  - High fraction of zero-day vulnerabilities
- Program Instability
- Difficult Debugging and Maintenance



In principle, yes. In practice...



Work to make C/C++ safe includes

- Static analysis (Clang Static Analyzer, Coverity, CppCheck, ...)
- Exploit mitigations (ASLR, CFI, guard pages, allocator hardening, ...)
- Sanitizers (ASan, UBSan, MemorySanitizer, ThreadSanitizer, GWPAsan, ...)
- Hardware-assisted techniques (MTE, PAC, PKRU, MPX, CET, CHERI, ...)
- Modern C++ (std::unique\_ptr, std::shared\_ptr, std::weak\_ptr, ...)
- Language dialects (Cyclone, CheckedC, CCured)



Work to make C/C++ sat performance

- Static analysis (Clang Static /
- Exploit mitigations (ASLR,
- Sanitizers (ASan, UBSar
- Hardware-assisted teg
- Modern C++ (std::up

tizer, ThreadSanitizer, GWPAsan, ...) C, PKRU, MPX, CET, CHERI, ...)

pages, allocator hardening, ...)

zer, Coverity, CppCheck, ...)

tr, std::weak\_ptr, ...)

• Lance Compatibility<sup>3, Checke</sup> security



Work to make C/C++ safe includes

- Static analysis (Clang Static And CppCheck, ...) Exploit mitigations (And Check And Ch





#### https://msrc.microsoft.com/blog/2019/07/we-need-a-safer-systems-programming-language



- Google: 90% of Android vulnerabilities are memory safety issues
- Apple: 60-70% of vulnerabilities in iOS and macOS are memory safety vulnerabilities
- Alex Gaynor: more than 80% of the exploited zero-day vulnerabilities were memory safety issues

https://security.googleblog.com/2019/05/queue-hardening-enhancements.htm https://langui.sh/2019/07/23/apple-memory-safety/ https://twitter.com/LazyFishBarrel/status/1129000965741404160



https://security.googleblog.com/2019/05/queue-hardening-enhancements.htm https://langui.sh/2019/07/23/apple-memory-safety/ https://twitter.com/LazyFishBarrel/status/1129000965741404160



https://chromium.googlesource.com/chromium/src/+/master/docs/security/rule-of-2.md



https://chromium.googlesource.com/chromium/src/+/master/docs/security/rule-of-2.md





- Expertise in source & target language
- Resources available to do the porting
- Resources available to maintain port



- Software isolation overheads tolerable
- Hardware isolation feasible
- Cannot support original and rewritten versions of same component





► ISOLATE

- Buy-in from current developers
- Expertise in source & target language
- from modern languages

- Software isolation overheads tolerable
- Hardware isolation feasible

meaningful security<sup>itten</sup> benefits in less time





→ ISOLATE

- Decide on a memory safe language
  - Chromium's strategy
    - Java for Android-specific code
    - Swift for iOS-specific code
    - Rust for third-party use
    - JavaScript or WebAssembly
  - C# good option for Microsoft-specific code
  - Swift for Apple-specific code



Step 1: Decide what to migrate

MIGRATE -→ ISOLATE

Step 1: Decide what to migrate



https://www.usenix.org/system/files/sec22-alexopoulos.pdf

### When to migrate and when to isolate?

How Long Do Vulnerabilities Live in the Code? A Large-Scale Empirical Measurement Study on FOSS Vulnerability Lifetimes

Nikolaos Alexopoulos, Manuel Brack, Jan Philipp Wagner, Tim Grube and Max Mühlhäuser Telecooperation Lab, Technical University of Darmstadt, Germany

#### Abstract

How long do vulnerabilities live in the repositories of large, evolving projects? Although the question has been identified as an interesting problem by the software community in online forums, it has not been investigated yet in adequate depth and scale, since the process of identifying the exact point in time when a vulnerability was introduced is particularly cumbersome. In this paper, we provide an automatic approach for accurately estimating how long vulnerabilities remain in the code (their *lifetimes*). Our method relies on the observation that while it is difficult to pinpoint the exact point of Reducing the number of vulnerabilities in software by finding existing ones and avoiding the introduction of new ones (e.g. by employing secure coding practices or formal verification techniques) is one of the primary pursuits of computer security.

Measurement studies on the different stages of the vulnerability lifecycle play an important role in this pursuit, as they help us better understand the impact of security efforts and improve software security practices and workflows. The community has produced a number of such outputs in recent years [5,8, 14, 18, 23, 28, 36]. A vulnerability's lifecycle, or







https://security.googleblog.com/2024/09/eliminating-memory-safety-vulnerabilities-Android.html



#### Total Memory safe and Memory Unsafe Lines of Code in AOSP

📒 Memory safe 📒 Memory Unsafe



Number of Memory Safety Vulns per Year



https://security.googleblog.com/2024/09/eliminating-memory-safety-vulnerabilities-Android.html

Gradual migration  $\rightarrow$  language interoperability

- Different memory management models
- Different type systems
- Different exception handling mechanisms
- Different data structure layouts and representations
- Different meta-programming facilities
- Different build systems

Gradual migration  $\rightarrow$  language interoperability

- Interoperability with C is doable
  - $\circ \quad \text{Simple type system} \leftrightarrow \text{complex type system}$
  - bindgen and cbindgen for Rust
- Interoperability with C++ is hard
  - $\circ \quad \text{Complex type system} \leftrightarrow \text{complex type system}$
  - $\circ$  Cxx, autocxx, Crubit for Rust



#### How to migrate a particular module/feature/library?

#### MANUALLY

- Best when existing code is poorly structured or lacks sufficient tests
- Not all features need to be ported

### UNSAFE TRANSLATION AUTOMATICALLY

- Best when existing code is high quality and tests have high code coverage
- Drop-in compatibility desired
- Need to keep up with upstream

 $\equiv$ 

#### How to migr

#### MAN

- Best when exis structured or la
- Not all features

The First Stable Release of a Memory Safe sudo Implementation

Josh Aas Aug 29, 2023

Prossimo is pleased to announce the <u>first stable release</u> of <u>sudo-rs</u>, our Rust rewrite of the critical sudo utility.

The sudo utility is one of the most common ways for engineers to cross the privacy boundary between user and administrative accounts in the ubiquitous Linux operating system. As such, its security is of the utmost importance.

The <u>sudo-rs project</u> improves on the security of the original sudo by:

- Using a memory safe language (Rust), as it's estimated that one out of three security bugs in the original sudo have been memory management issues
- Leaving out less commonly used features so as to reduce attack surface
- Developing an extensive test suite which even managed to <u>find bugs</u> in the original sudo

CALLY

code is high quality code coverage y desired th upstream

#### How to migr

#### MAN

💋 PROSSIMO

- Best when exis structured or la
- Not all features

#### Porting C to Rust for a Fast and Safe AV1 Media Decoder

Stephen Crane Sep 9, 2024

AV1 is an increasingly important video format and it needs a memory safe, high performance decoder. We worked with the team at <u>Immunant</u> to develop **rav1d**, a Rust-based port of **dav1d**, a C decoder. This is the first of two blog posts about how the team approached this effort.

— Josh Aas, Head of ISRG's Prossimo project

Complex data parsing is one of the most security-critical operations in modern software. Browsers must decode untrusted audio and video inputs encoded with extremely complicated formats in real time. Memory safety bugs in this decoding process are disastrous and common. For example, researchers fuzzing H.264 decoder implementations have demonstrated that these decoders are a <u>dangerous source of bugs</u>. AV1 is a similarly complex, widely used video format. We need a <u>memory safe, performant implementation</u> of AV1 format parsing to avoid parsing vulnerabilities in heavily targeted software such as browsers.

To create this fast and safe AV1 decoder, we have ported an existing high performance AV1 decoding library, **dav1d**, from C to Rust: **rav1d**. Our implementation is drop-in compatible with the **dav1d** C API. Format

 $\equiv$ 

#### CALLY

code is high quality code coverage y desired th upstream

#### Quick aside on Internet Security Research Group

Funders

Josh Aas (Co-Founder)

Sarah Gran (VP of Comms)



cisco.

aws

Cacton family giving

🙎 Fly.io

Sovereign Tech Fund



craig newmark philanthropies







🗿 shopify

https://www.memorysafety.org/ and https://isrg.org/

#### Porting dav1d from C to Rust using c2rust

Stage 0 Stage 1 .c .rs Unsafe C Transpiler Unsafe Rust Stage N Stage N+1 Rewriting or .rs .rs Intermediate Rust 돴 Improved Rust Refactoring iterate



#### C2Rust transpiler example inputs and outputs

#### C source code 1 \* void insertion\_sort(int const n, int \* const p) { 2 3 for (int i = 1; i < n; i++) { 4 int const tmp = p[i]; 5 int j = i; 6 while (j > 0 && p[j-1] > tmp) { 7 p[j] = p[j-1];8 j--; 9 3 10 p[j] = tmp;11 12 } Upload insertion V Load Sample Translate

#### Generated Rust source code #![allow(dead\_code, mutable\_transmutes, non\_camel\_case\_types, non\_snake\_case, non 1 2 #[no\_mangle] 3 - pub unsafe extern "C" fn insertion\_sort(n: libc::c\_int, p: \*mut libc::c\_int) { let mut i: libc::c\_int = 1 as libc::c\_int; 4 5 while i < n { 6 let tmp: libc::c\_int = \*p.offset(i as isize); let mut j: libc::c\_int = i; 7 8 while j > 0 as libc::c\_int && \*p.offset((j - 1 as libc::c\_int) as isize) 9 \*p.offset(j as isize) = \*p.offset((j - 1 as libc::c\_int) as isize); 10 j -= 1; 11 j; 12 3 13 \*p.offset(j as isize) = tmp; 14 i += 1; 15 i: 16 3 17 18 Download output.rs

Try it out at https://c2rust.com or https://github.com/immunant/c2rust

#### C2Rust transpiler example inputs and outputs



Try it out at <a href="https://c2rust.com">https://github.com/immunant/c2rust</a>

#### Porting dav1d from C to Rust using c2rust

- Initial transpile took almost no time (~1 day)
  - Transpilation was beneficial w.r.t. testing & tracking upstream changes
  - $\circ$   $\,$   $\,$  Transpiler output was ~5x as verbose as input C code
- ~90% time spent "cleaning up" the transpiler output (~12mos)
  - Full test suite run on each commit to catch regressions early
  - Extensive use of C preprocessor macros presented a challenge
  - Threading patterns difficult to translate to Rust efficiently
- ~10% time spent optimizing, handling language differences, CI config, etc.
  - $\circ$  Difficult to track down where we lost performance vs. the C version
    - Currently seeing a ~6% overhead on x86\_64
  - Details <u>https://www.memorysafety.org/blog/rav1d-performance-optimization/</u>

#### Porting dav1d from C to Rust using c2rust



#### Lifting C code to Safe Rust

- Use static analysis to identify lifting opportunities
- Use dynamic analysis to augment static analysis
- Combine results in to pointer derivation graph

immunant × |galois|

- Assign permissions to pointers
- Rewriting engine maps C types to safe Rust types



#### Lifting C code to Safe Rust

- Use static analysis to identify lifting opportunities
- Use dynamic analysis to augment static analysis
- Combine results in to pointer derivation graph

immunant × |galois|

- Assign permissions to pointers
- Rewriting engine maps C types to safe Rust types

Write	Unique	Free	Offset	Ptr type
				&т
1	1			&mut T
1				Cell <t></t>
			1	&[T]
1	1		1	&mut [T]
✓ / X	1	1		Box
✓ / X	1	1	1	Box<[T]>

#### **Heap Allocations**

```
pub struct buffer {
    // was: *mut i8
    pub ptr: Option<DynOwned<Box<[i8]>>>,
    pub used: u32,
    pub size: u32,
}
```

(DynOwned<T> tracks ownership dynamically, like Option<T>)

#### immunant x |galois|

#### **Heap Allocations**



#### immunant × |galois|

#### DARPA TRACTOR



https://www.darpa.mil/program/translating-all-c-to-rust



#### Current state of the art (not all translations are equally good)





- Hardware-based isolation
  - CHERI
  - TrustZone
  - Intel SGX
  - AMD SVM
  - (Apple Secure Enclave)





- Virtualization
  - Android Virtualization Framework
    - https://source.android.com/docs/core/virtualization/architecture
  - Extends Linux KVM with protected VM feature
  - Supports both ARM and x86\_64

#### When to migrate and when to isolate? VIALSOSOLATE BUTALSOSOLATE **MIGRATE** AVYNEIGHT AND Virtualization HEALY COMPLEX zation/architecture Android Virtuali--Exter Suppo -



- Process level-sandboxing
  - Separate high and low privilege components in separate OS processes
  - Used by most modern browsers
  - Highly portable and mature isolation strategy
    - Non-trivial to apply to monolithic applications

When

https://www.usenix.org/conference/enigma2021/presentation/palmer

#### THE LIMITS OF SANDBOXING AND NEXT Proc STEPS

Note: Presentation times are in Pacific Standard Time (PST).

Wednesday, February 03, 2021 - 9:50 am-10:20 am

Chris Palmer, Google Chrome Security

#### Abstract:

Privilege separation and reduction ("sandboxing") has significantly improved software security, and in many applications is a baseline requirement for safe design. (In fact, there are still many applications that can and should adopt sandboxing.)

Although necessary, sandboxing is not sufficient by itself. The designs and implementations of real-world operating systems put a ceiling on the effectiveness and applicability of sandboxing. From years of experience shipping Chromium, we have learned that (1) Chromium is at or near the limit of how much safety it can practically provide with privilege separation and reduction; and (2) we still need to provide greater resilience.

Therefore, we must find and develop additional security mechanisms. Our primary approach is now working toward increased memory safety. Where sandboxing limits the value attackers gain from exploiting vulnerabilities, memory-safe(r) code can eliminate vulnerabilities altogether or make it infeasible to use them in an exploit chain.

ses

When

https://www.usenix.org/conference/enigma2021/presentation/palmer

applicability of sandboxing. From years of experience shipping Chromium, we have learned that (1) Chromium is at or near the limit of how much safety it can practically provide with privilege separation and reduction; and (2) we still need to provide greater resilience.

Therefore, we must find and develop additional security mechanisms. Our primary approach is now working toward increased memory safety. Where sandboxing limits the value attackers gain from exploiting vulnerabilities, memory-safe(r) code can eliminate vulnerabilities altogether or make it infeasible to use them in an exploit chain.

shipping Chromium, we have learned that (1) Chromium is at or near the limit of how much safety it can practically provide with privilege separation and reduction; and (2) we still need to provide greater resilience.

Therefore, we must find and develop additional security mechanisms. Our primary approach is now working toward increased memory safety. Where sandboxing limits the value attackers gain from exploiting vulnerabilities, memory-safe(r) code can eliminate vulnerabilities altogether or make it infeasible to use them in an exploit chain.

- Library-level sandboxing
  - Avoids the need to split application into multiple processes
- WebAssembly
- RLBox
  - https://rlbox.dev/
  - Can use WebAssembly or PKRU as backends
- Hardware-assisted library sandboxing
  - https://github.com/immunant/IA2-Phase2/

#### Takeaway

- Memory safe languages delivers assurance that previous practices lacks.
- Follow the rule of two!

(don't process untrusted input with unsafe code)

• We need to dramatically **lower the cost of migrating code** to memory safe languages.





#### Memory Safety

- Means that software or a programming language is designed to catch memory bugs
- Memory safe languages automate memory management only allow safe memory accesses
- Most languages are memory safe
- C/C++ are the only languages that are still widely used but not memory safe
  - Prizes efficiency and flexibility, programmer is in full control and has full responsibility
  - Routinely used to write low-level, highly-privileged code
- Consequences of Memory Unsafety
  - Security vulnerabilities
  - Program instability
  - Difficult debugging and maintenance
- Further reading
  - <u>https://www.internetsociety.org/resources/doc/2023/how-to-talk-to-your-manager-about-memory-safety</u>
  - https://www.memorysafety.org/docs/memory-safety/
  - https://storage.googleapis.com/gweb-research2023-media/pubtools/7665.pdf

#### **Further Reading**

- <u>https://www.internetsociety.org/resources/doc/2023/how-to-talk-to-your-manager-about-memory-safety</u>
- https://www.memorysafety.org/docs/memory-safety/
- https://storage.googleapis.com/gweb-research2023-media/pubtools/7665.pdf