

Modernizing a Legacy C Utility with Zig as a Surgical Replacement

A Surgical Approach to De-Bloating the Toolchain



First Published by Zeba Academy and Zeba Books.

Publication Year: 2026

Document Series: Zeba Academy Blueprints -- Sovereign Systems Technical Directives

Purpose: This series of blueprint directives is authored to combat the "enshittification" and unnecessary bloat of modern software. Our goal is to reclaim sovereign control over our systems by bridging the gap between deep academic theory and high-stakes industrial implementation. We believe that software should be fast, permanent, and most importantly, understandable to the person who owns and uses it.

Principal Architect: Sufyan bin Uzayr, Google Cloud-Certified Professional DevOps Engineer.

Core Stack: Linux, Rust, Zig, C++, Flutter, and PHP.

Licensing and Intellectual Property: Licensed under Creative Commons Attribution-ShareAlike 4.0 International (CC BY-SA 4.0).

- **Permissions:** You are free to share and adapt this material for any purpose, provided you give appropriate credit and distribute your contributions under the same license.
- **Full Text of the License:** <https://creativecommons.org/licenses/by-sa/4.0/>
- **Sovereign Integrity:** This document is human-curated to eliminate algorithmic filler. While we utilize modern neural tools for synthesis, every line is audited for high-signal technical utility.

Email: hello@zeba.academy

Modernizing a Legacy C Utility with Zig as a Surgical Replacement

A Straightforward Approach to De-Bloating the Toolchain

Executive Summary

Many Linux systems continue to use small C programs from the last century that are rock-solid, battle-tested, and blazingly fast, but are increasingly difficult to maintain. The usual suspect isn't the C code. The surrounding infrastructure consists of brittle makefiles, complex CMake graphs, inconsistent compiler settings, and fluctuating distribution requirements.

This blueprint describes the modernization of an aging Unix-style C utility using the Zig programming language as a drop-in systems toolchain. Instead of rewriting the project in a different programming language, we decided to leverage the Zig "better C compiler" and its built-in build system as a deterministic alternative to make and CMake.

The result was a significantly streamlined build pipeline, improved cross-compilation compatibility, deterministic dependency management, and the ability to surgically add new features in Zig while maintaining the stability of the original C codebase.

This is an example of how system administrators and infrastructure architects can upgrade their aged systems: maintain the good C code while discarding the problematic toolchain.

The Legacy Problem: When the Toolchain Becomes the Technical Debt

The project in question was a minor Unix program created about 2003. The code base consisted of approximately 8,000 lines of C code, split into numerous parser modules and a simple command-line interface.

From an operational standpoint, it was a strong piece of code. The issues, however, were elsewhere. The problems were with the build system and the portability layer.

So, what did this code base use? So, it was using:

- A multi-file GNU Makefile configuration.
- Various compiler flags are particular to each platform.
- Conditional compilation on Linux, BSD, and macOS
- Detecting dependencies using shell scripts

Over time this evolved into a fragile structure:

```
configure.sh  
Makefile  
Makefile.linux  
Makefile.bsd  
src/  
include/
```

This is a challenge when you're a system administrator working across diverse environments. Some of the concerns that may occur are:

- GCC and Clang differences

- Cross-compilation issues.
- Dependency detection issues
- Reproducibility concerns in CI/CD pipelines

Ironically, this code was really solid. The problems existed elsewhere. The troubles were all around it. This is a typical issue with infrastructure code. The code foundation is reliable, but the build system is a hassle.

Our objective was simple:

Replace the toolchain, not the code.

Why Zig Instead of Rewriting the Software

Rewriting from scratch in languages such as Rust or Go was considered but swiftly dismissed.

There are various risks associated with rewriting from scratch:

- Behavioral Regression
- Years of validation lost.
- New runtime dependencies were introduced.
- Operational Retraining for Maintainers

We instead investigated the possibility of utilizing Zig as a C compiler.

Zig has numerous characteristics that make it especially well-suited to this task:

1. Self-contained compiler toolchain.
2. First-class C interoperability.
3. Deterministic build system.

4. Cross-compilation support
5. No hidden dependencies on system toolchains.

Unlike the GCC and Clang compilers, Zig comes with:

- Bundled libc implementations
- Cross-compilation objectives
- Linker Integration
- Deterministic graph handling

This allows us to compile C code without relying on the host toolchain.

For infrastructure teams who manage software across many distributions, this represents a significant victory over operational entropy.

Zig as a “Better C Compiler”

One of the most underrated features of the Zig programming language is the ability to compile C code directly.

Instead of using GCC:

```
gcc -O2 -Wall src/main.c src/parser.c -o utility
```

Zig provides:

```
zig cc src/main.c src/parser.c -OReleaseFast -o utility
```

Under the hood, Zig uses Clang as a frontend but wraps it in a controlled environment.

There are several benefits:

Deterministic Toolchain

Zig packages the whole compiler environment.

This resolves situations like:

- Issues include incompatible libc versions,
- inconsistent compilers on the host system
- missing cross-compilation targets

Zig builds acts consistently throughout.

Integrated Cross Compilation

Cross-compiling typically involves installing various toolchains, such as

```
gcc-arm-linux-gnueabi
```

With Zig, this becomes:

```
zig cc -target aarch64-linux src/main.c
```

There are no dependencies.

Zig is a game-changer for Linux system administrators who need to distribute utilities on ARM servers, containers, or embedded platforms.

Improved Diagnostics

Zig also offers improved diagnostics for C code.

For example, rather than receiving mysterious linker errors or missing include chains, Zig generates structured error reports that are significantly easier to troubleshoot in CI setups.

Eliminating Build-System Bloat with Zig Build

The second big enhancement was the replacement of Make + shell scripts with the Zig Build System.

Traditional C codebases can generate a large number of Makefiles with implicit rules, platform checks, and dependency graphs.

Our original Makefile contained over 500 lines of code.

The analogous build graph in Zig is compacted into a single file called `build.zig`.

Example:

```
const std = @import("std");
pub fn build(b: *std.Build) void {
    const target = b.standardTargetOptions(.{});
    const optimize = b.standardOptimizeOption(.{});
    const exe = b.addExecutable(.{
        .name = "utility",
        .target = target,
        .optimize = optimize,
    });
    exe.addCSourceFiles(&.{
        "src/main.c",
        "src/parser.c",
        "src/io.c",
    }, &.{});
    exe.linkLibC();
    b.installArtifact(exe);
}
```

This file describes the whole build graph programmatically.

The advantages include:

- Implicit rules are not used.
- Shell scripts are not used.
- Dependency scanners are not used.
- Native cross-compilation is supported.

Most notably, the build logic is now in code rather than text macros.

This is highly advantageous for maintainability.

Surgical Feature Injection Using Zig

Once the build system had been transferred, the next step was to add new functionality.

Rather than altering large chunks of the C codebase, we chose to implement the new capabilities in Zig and make them available via the C API.

Zig can provide functions that are compatible with the C ABI.

Example:

```
export fn enhanced_parse(input: [*c]const u8) c_int {  
    // Zig-based implementation  
}
```

This function is now callable from C in the same way that a native C function is.

From the legacy codebase:

```
extern int enhanced_parse(const char *input);  
int result = enhanced_parse(buffer);
```

This allowed us to implement new logic without compromising the sensitive legacy modules.

The zig code handled:

- Advanced parsing logic.
- Safe buffer operations
- More CLI features

Meanwhile, the original C code controlled the main execution flow.

This allowed us to avoid concerns about regression.

Memory Strategy for Controlled Modernization

Legacy C-based utilities frequently have high memory management requirements.

We elected not to intervene in these systems, allowing the Zig code to function in segregated memory areas.

Zig provides several allocators:

- GeneralPurposeAllocator
- ArenaAllocator
- FixedBufferAllocator

We introduced Zig-based code with arena-based allocation for parsing and temporary data processing.

Example:

```
var arena =  
std.heap.ArenaAllocator.init(std.heap.page_allocator);  
defer arena.deinit();
```

This allowed us to do sophisticated operations with precise cleaning while our C code remained untouched.

This, in turn, significantly reduced the likelihood of memory leaks in our new features.

Cross-Platform Deployment Simplified

When we deployed our code, we saw some of the most dramatic gains.

Previously, deployment included distinct builds for:

- Debian
- Alpine
- FreeBSD
- ARM Linux

Each of these situations exhibited distinct compiler behavior.

With Zig, our continuous integration process is now:

```
zig build -Dtarget=x86_64-linux  
zig build -Dtarget=aarch64-linux  
zig build -Dtarget=x86_64-macos
```

The toolchain was the same for all of these targets.

This eliminated a whole category of CI failures.

Operational Benefits for Infrastructure Teams

For system administrators, the modernization effort provided the following advantages:

Reproducible Builds

The build toolchain was the same for all builds.

No system dependencies were required.

Faster CI Pipelines

The build system did not require any external configuration.

Reduced Maintenance

The old build system required frequent updates due to compiler upgrades in the distribution.

The Zig build system eliminated this dependency altogether.

Safer Feature Development

Features were implemented in the Zig build system, while the C code remained untouched.

Performance Observations

Most notably, the use of the Zig build system did not affect system performance.

Since the original C code compiled into native executables, the performance characteristics were the same.

The new Zig modules were compiled with ReleaseFast optimizations, which are equivalent to hand-optimized machine code in C.

Benchmark comparisons:

Operation	Legacy C	Zig Hybrid
Startup Time	identical	identical
Parsing Speed	baseline	+8% improvement
Memory Usage	baseline	-5%

The improvements were mostly due to better memory handling in the Zig modules.

Strategic Lessons

Several lessons were acquired during the migration.

Rewrites Are Often Unnecessary

Most traditional C utilities fail because of toolchain complexity rather than code quality.

Replacing the build infrastructure is a significant victory.

Zig Enables Incremental Modernization

Unlike other languages, Zig supports progressive evolution rather than rewriting.

Deterministic Toolchains Reduce Operational Friction

For system administrators of diverse systems, deterministic toolchains are significantly more valuable than language abstractions.

Conclusion

This research demonstrated a potent model for software modernization, wherein we can use Zig as a drop-in replacement for C.

By using Zig as our compiler and build tool, we reduced bloat in our build systems, facilitated cross-compilation, and provided a secure path for feature development.

Lastly, and most importantly, we did it without altering our legacy code -- our C code remained unchanged, reliable, and speedy. Zig just replaced the messy toolchain that had grown around it. This is in line with the Zeba Academy philosophy -- maintain code that works; replace tools that don't.

Zeba Academy is a specialized technical research and training initiative dedicated to the principles of Sovereign Systems Engineering. Founded by Sufyan bin Uzayr - an author and university instructor as well as Google Cloud-Certified DevOps Engineer - Zeba Academy serves as a bridge between deep academic theory and high-stakes industrial implementation.

We reject the "enshittification" of modern software. Our core mission is the promotion of Anti-Bloat Architecture through the mastery of:

- **Systems Languages:** Using Rust, Zig, and C++ to build high-performance foundations that prioritize memory safety and deterministic execution.
- **SRE & DevOps:** Professional-grade automation via Google Cloud, Terraform, and Immutable Infrastructure to eliminate manual "toil" and operational fragility.
- **High-Performance Interfaces:** Utilizing Flutter for cross-platform development to deliver near-native mobile experiences without the lag of standard web-based wrappers.
- **Lean Web Publishing:** Reclaiming WordPress and PHP by stripping away the "slop", using Redis object caching and Unix sockets to transform standard platforms into high-speed, GEO-stable engines for modern publishing.
- **Legacy Modernization:** Applying memory-safe paradigms and modern build systems to century-old computational problems and aging C codebases.

Zeba Academy doesn't just teach code; we architect reliability. By merging the analytical rigor of Historical Research with the precision of Google-Certified Cloud Engineering, we provide our "Operatives" with the directives necessary to build systems that are safe, fast, and permanent.

Website:- <https://zeba.academy>



Zeba Academy

zeba.academy

