FRAMEWORK-FREE SOFTWARE ENGINEERING

BUILDING ROBUST SYSTEMS WITHOUT HEAVYWEIGHT DEPENDENCIES

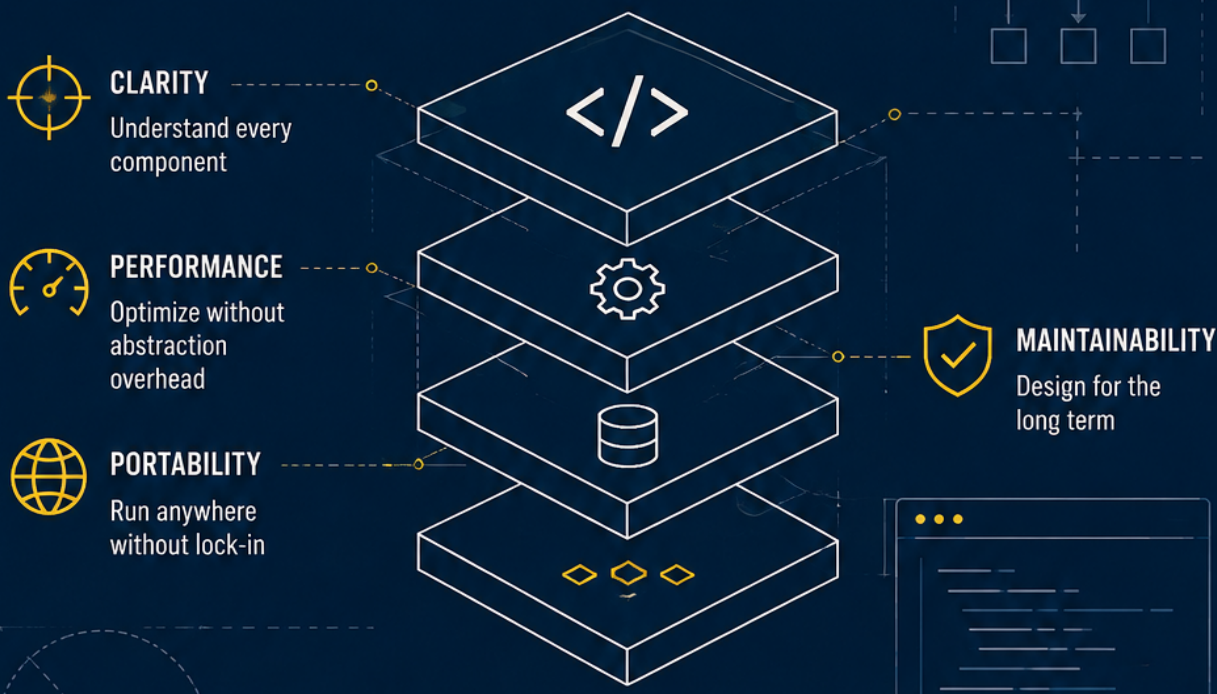

WILLIAM HARRINGTON

Framework-Free Software Engineering

Building Robust Systems Without Heavyweight Dependencies

Steve Publications

This book is available at
https://leanpub.com/framework-freesoftwareengineering

This version was published on 2026-08-01

This is a Leanpub book. Leanpub empowers authors and publishers with the Lean Publishing process. Lean Publishing is the act of publishing an in-progress ebook using lightweight tools and many iterations to get reader feedback, pivot until you have the right book and build traction once you do.

© 2026 Steve Publications

Contents

Building Robust Systems Without Heavyweight Dependencies

This book shows you how to design, build, and operate production-grade software systems without relying on heavyweight frameworks. Through architectural principles, practical patterns, complete code examples, and real-world case studies, you will learn to make deliberate engineering decisions that maximize clarity, performance, portability, and long-term maintainability. Whether you are evaluating a move away from framework-heavy stacks or simply want deeper understanding of what your tools hide, this book provides the knowledge to build software on your own terms.

Introduction: The Framework Question

A senior engineer at a fast-growing startup walked into my office in 2019 and said something I have never forgotten. We had just spent three weeks debugging a memory leak that turned out to live inside a popular web framework we used for our main API service. After reading through the framework source code, tracing garbage collector pauses, and correlating heap dumps with request patterns, we finally found it: a subtle bug in the framework's built-in middleware that created circular references under specific concurrency conditions. The fix required patching a dependency nobody on the team fully understood, submitting a pull request upstream, and waiting two weeks for the maintainers to review and release it.

"We could have written this whole service ourselves by now," he said. "And we would have known exactly where every byte of memory was going."

He was not wrong. But he was also not telling the whole story. Those same framework features had let us ship our initial API in three days instead of three weeks. They had given us routing, request parsing, middleware composition, and JSON serialization out of the box, while we focused on business logic. Without them, we would have taken much longer to get to market. The framework paid for itself many times over during those first months.

The problem was not the framework itself. It was that we stopped understanding what it did. We accepted its abstractions as given. When things worked, they worked magically. When they broke, we were helpless until someone with deep framework knowledge could dig through layers of indirection to find the root cause.

This book is about making sure you are never in that position again. It is not an argument against frameworks. Frameworks solve real problems for real teams. But it is an argument for understanding, control, and deliberate choice. If you can build something without a framework, you understand what a framework gives you when you use one. You know the trade-offs. You can spot when abstraction has become a liability instead of an asset. And most importantly, you have the skills to walk away when a framework stops serving you.

What This Book Will Do

After reading this book, you will be able to:

- Design and implement clean, maintainable architectures without framework-enforced structure
- Build production-ready HTTP APIs, microservices, and background workers using only language primitives and carefully chosen libraries
- Implement dependency injection, configuration management, logging, error handling, testing, security, observability, and deployment pipelines from first principles
- Evaluate frameworks critically and choose tools based on actual requirements rather than popularity or convention
- Migrate existing framework-heavy systems toward simpler, more controllable architectures when the cost of complexity outweighs the benefit of convenience

Who This Book Is For

This book assumes you are already a competent software developer. You have built applications before. You have used at least one mainstream framework and understand basic concepts like routing, middleware, dependency injection, and database access. If you are a senior engineer, tech lead, or architect looking to deepen your understanding of what lies beneath the tools you use every day, this book is for you.

If you are new to software development, this book will be dense. It covers topics that experienced engineers debate: architectural patterns, concurrency models, security protocols, deployment strategies. You will benefit more from it after you have spent some time building real systems and encountering the problems these patterns exist to solve.

A Note on Language Choices

The examples in this book use Go as the primary language. Go was chosen deliberately: its standard library provides enough building blocks to build

production services without external frameworks, its concurrency model is elegant and practical, and its culture values simplicity over ceremony. Secondary examples appear in Rust and Node.js where they illuminate different points or demonstrate alternative approaches. The principles apply across languages, even if the syntax differs.

How to Read This Book

The chapters build on each other logically. Chapter 1 defines what framework-free development actually means and motivates why it matters. Chapter 2 provides historical context for how we arrived at today's framework-heavy landscape. Chapters 3 through 6 establish the architectural foundations: SOLID principles, clean architecture, modular design, dependency management, configuration, and bootstrapping. Chapters 7 through 11 cover specific implementation domains: API design, error handling and observability, persistence without ORMs, concurrency, and security. Chapter 12 covers testing strategies. Chapter 13 covers build, deployment, and operations. Chapter 14 presents real-world case studies of systems built using these principles. Chapter 15 synthesizes everything into a practical decision framework for choosing between framework-based and framework-free approaches in your own projects.

You can read the book linearly from start to finish, or you can jump to chapters relevant to your immediate needs. If you are already comfortable with SOLID principles and clean architecture, you can skim Chapters 3 and 4. If your main interest is API design, Chapter 7 stands on its own. But for the full picture, reading sequentially will give you the most coherent understanding of how all these pieces fit together.

Chapter 1: The Case for Framework-Free Development

What Framework-Free Actually Means

The phrase "framework-free" invites misunderstanding. It does not mean you write everything from scratch, including your HTTP server, JSON parser, and cryptographic primitives. That would be foolish reinvention of the wheel. It also does not mean you reject all structure, convention, or opinionated design. Good software needs constraints to keep complexity manageable.

What framework-free means is a specific distinction between frameworks and libraries, combined with a deliberate preference for libraries over frameworks wherever possible. The key difference lies in who controls the flow of execution.

A library is a collection of functions or classes you call from your own code. You remain in control. Your application invokes library routines when it needs them, and execution returns to your code afterward. This is sometimes called the Hollywood Principle in reverse: you call out, rather than waiting to be called. Standard library packages are the purest form of this pattern. When you import a JSON serialization package and call its marshal function, you decide when, where, and with what data it runs.

A framework inverts this control. Instead of you calling the framework, the framework calls your code. You write handlers, callbacks, or subclasses that plug into a larger structure managed by the framework. The framework owns the main loop, the request lifecycle, the dependency graph. Your code becomes a set of extensions to someone else's architecture. This inversion of control is what distinguishes a framework from a library, and it carries consequences that are not always obvious when you first adopt one.

Framework-free development means building your applications primarily with libraries while maintaining your own control over the execution flow, the dependency graph, the lifecycle management, and the architectural boundaries. You wire your own components together. You decide how requests are

routed, how errors propagate, how configuration is loaded, and how services start and stop. You may use opinionated tools for specific tasks, but you refuse to surrender overall control to a single monolithic framework that dictates your application structure.

Consider three concrete examples across different ecosystems:

In Node.js, using Express or Fastify means your application conforms to their middleware chain model. Routes are registered with the framework object. Request and response objects are framework-managed wrappers around the underlying HTTP streams. Middleware executes in a predetermined order controlled by the framework's router. This is a framework. Using the built-in http module directly, writing your own request parser, routing logic, and middleware composition, gives you full control. You still use libraries for JSON parsing or logging, but you own the architecture.

In Java, Spring Boot defines your application structure through annotations, dependency injection containers, auto-configuration mechanisms, and lifecycle callbacks. Even a simple REST endpoint requires understanding Spring's component scanning, bean lifecycle, MVC dispatcher, and configuration property binding. This is a heavyweight framework. Writing a plain Java application that uses the built-in HTTP server or a lightweight library like Netty directly, with manual dependency wiring and explicit configuration loading, is framework-free.

In Go, the situation is different because Go's culture already leans toward minimalism. The standard library's net/http package provides a fully functional HTTP server without requiring any external framework. Adding a router like chi or mux gives you more convenient routing without inverting control: you still own the main function, the handler registration, and the application lifecycle. This is why Go services are often framework-free by default, even when they use third-party libraries extensively.

The goal is not to eliminate all frameworks from your toolkit. Frameworks excel at specific domains where their opinions match your needs perfectly. React has earned its place in frontend development because its component model and virtual DOM solve real problems that would be painful to reimplement. Django remains popular for content-heavy web applications because its admin interface, ORM, and authentication system provide enormous value for standard use cases. The question is whether the framework's structure serves your application or whether your application contorts itself to fit the

framework's expectations.

The Cost of Abstraction Layers

Every abstraction has a cost. This is not a judgment against abstractions themselves; they are essential tools for managing complexity. The statement is simply that costs exist, they accumulate, and they are often invisible until it is too late to avoid them easily.

Joel Spolsky formulated the Law of Leaky Abstractions: all non-trivial abstractions leak to some degree [1]. Underlying details eventually surface when the abstraction cannot fully hide complexity. A database connection pool abstracts away individual connections, but you still need to understand timeouts, stale connections, and connection limits when your application misbehaves under load. An ORM abstracts SQL queries, but you still need to understand joins, indexes, and query plans when performance degrades. The abstraction saves you time working in the normal case, but it does not save you time learning what happens in the edge cases [2].

Abstraction costs fall into several categories:

Execution cost refers to the computational overhead introduced by indirection. Each additional layer adds function calls, object allocations, serialization steps, and runtime interpretation. Butler Lampson warned that if there are six levels of abstraction and each costs fifty percent more than reasonable, the performance will suffer dramatically [3]. In many applications this overhead is negligible. A single extra function call per request adds microseconds to latency. But in high-throughput systems handling millions of requests per second, those microseconds compound into seconds of wasted compute time and significant infrastructure cost.

Mental cost refers to the cognitive burden of understanding multiple layers of indirection. When debugging a production issue, you must trace through framework code you did not write, understand its internal state machines, interpret its error messages, and correlate its behavior with your own application logic. The deeper the abstraction stack, the harder this becomes. A framework that logs "request failed at stage 3 of pipeline" requires you to learn what stage 3 means in that framework's lifecycle before you can diagnose anything.

Control cost refers to the difficulty of fine-tuning behavior when the abstraction does not expose the knobs you need. Want to change how a

framework handles connection timeouts? Override a configuration property if one exists, subclass a protected method if that is allowed, or monkey-patch internal behavior as a last resort. None of these are clean solutions. The framework decided what was configurable when it was designed, and you are stuck with those decisions unless you abandon the framework entirely.

Dependency cost refers to the supply chain risk of relying on external projects. Every framework depends on other libraries, which depend on their own dependencies, creating a tree that may contain thousands of transitive packages. Each dependency is a potential source of bugs, security vulnerabilities, breaking changes, or abandoned maintenance. The Log4j vulnerability in 2021 affected millions of applications because they depended on a logging library they did not directly maintain or even fully understand [4]. Framework-free systems with minimal dependencies have a smaller attack surface and fewer points of failure.

Learning cost is the time required to master a framework's conventions, patterns, APIs, and quirks. New team members must learn not just your application domain but the framework's way of doing things before they can be productive. Frameworks also create lock-in: developers trained on Spring Boot may struggle to write plain Java services. Developers accustomed to Rails conventions may find framework-free development uncomfortable because nothing tells them where files should go or how components should interact. The framework becomes part of the team's mental model, and removing it requires retraining.

Daniel J. Sturtevant analyzed eight years of data from a large software firm in his 2013 MIT dissertation and found empirical proof that architectural complexity significantly impacts productivity, defect rates, and developer retention [5]. Complex abstractions make code harder to understand, which makes it harder to modify safely, which leads to more bugs and slower delivery. The team members who can navigate the complexity become bottlenecks. Those who cannot become frustrated and leave.

None of this means you should never use abstractions. It means you should use them deliberately, understanding what they cost and ensuring the benefit justifies the expense. Framework-free development is not about avoiding all structure; it is about choosing structure that is transparent, minimal, and under your control.

Control, Transparency, and Cognitive Load

The central argument for framework-free development is control over your system's behavior, structure, and evolution. When you build without a heavyweight framework, you make every architectural decision consciously rather than inheriting someone else's decisions by default.

Control means you understand the full lifecycle of a request as it moves through your application. You know exactly which code runs when the server starts, how incoming HTTP requests are parsed and routed, how business logic is invoked, how database queries are executed, how responses are serialized and sent back to the client, and how errors are handled at each step. There is no hidden middleware chain executing before your handler runs. There is no framework-owned event loop processing events in ways you cannot observe. The path from network socket to business logic and back is explicit and traceable.

This transparency has immediate practical benefits. When a production incident occurs, you can diagnose it faster because there are fewer unknown layers to investigate. You can profile performance more accurately because you know where time is being spent without having to account for framework overhead. You can explain system behavior to stakeholders, new team members, and on-call engineers because the architecture is visible in the code rather than buried inside a framework's internals.

Cognitive load is reduced when your application structure matches your mental model instead of a framework's conventions. Frameworks impose their own organization: Spring Boot expects packages organized by layer with configuration classes annotated in specific ways. Rails expects models, views, and controllers in predetermined directories with naming conventions that must be followed precisely for auto-discovery to work. These structures are reasonable for the use cases the frameworks were designed for, but they become burdensome when your application does not fit those patterns neatly.

Framework-free code can organize itself around your domain rather than a framework's architecture. You might structure a service by feature modules, each containing its own handlers, business logic, and data access code. You might use a layered architecture with clear boundaries between presentation, application, domain, and infrastructure concerns. You might adopt hexagonal architecture with ports and adapters isolating core logic from external depen-

dencies. The choice is yours because no framework is dictating the structure through conventions you must follow.

This freedom carries responsibility. Without a framework enforcing consistency, teams can drift into inconsistent patterns across different parts of the codebase. Some developers might organize code one way while others choose a different approach, leading to confusion and maintenance overhead. This is why framework-free development requires deliberate architectural decisions documented in style guides and enforced through code review. The structure must come from the team's shared understanding, not from a framework's auto-configuration.

When Frameworks Earn Their Keep

Framework-free development is not universally superior. Frameworks solve real problems, and there are situations where their benefits clearly outweigh their costs. Understanding when to embrace a framework is as important as understanding when to avoid one.

Frameworks excel at accelerating development for standard use cases. If you are building a content management system, an e-commerce platform, or a CRUD-heavy internal tool, frameworks like Django, Rails, or Laravel provide authentication, authorization, database migrations, admin interfaces, and form handling out of the box. Building all of that from scratch would take weeks or months of work that the framework delivers in hours. For these applications, the framework's opinionated structure is a feature rather than a limitation because the application fits the framework's expectations perfectly.

Frameworks also provide consistency across teams. Large organizations with many development teams benefit from a shared framework that standardizes how services are built, tested, deployed, and operated. New developers can move between projects without learning entirely new architectures. Code reviews are easier when everyone follows the same patterns. Operational tooling can assume a common structure for logging, metrics, health checks, and configuration. The cognitive cost of learning one framework is amortized across many projects and many team members.

Frameworks mature over time and accumulate battle-tested solutions to common problems. Spring Boot's ecosystem includes well-maintained modules for security, messaging, caching, distributed tracing, and cloud integration that have been tested by millions of users in production environments.

Reimplementing all of this functionality yourself would be expensive and risky. Using the framework gives you access to this collective knowledge without reinventing the wheel.

The key question is whether your application fits the framework's strengths or fights against them. If your use case aligns with what the framework was designed for, embrace it fully. If your application has unique requirements that the framework does not support well, evaluate whether the cost of working around those limitations exceeds the benefit of the features you do get. Sometimes the answer is to use the framework selectively: adopt it for areas where it adds value while building custom solutions for areas where it does not fit.

The rest of this book teaches you how to build robust systems without frameworks so that you can make this evaluation knowledgeably. You will learn the patterns and techniques that frameworks encapsulate, so you understand what you are getting when you use one and what you are giving up in terms of control, performance, and transparency. Armed with that understanding, you can choose the right tool for each job instead of defaulting to whatever framework is popular or familiar.

The framework question has no single correct answer. It depends on your application's requirements, your team's skills, your operational constraints, and your long-term vision for the system's evolution. But it should always be a conscious choice based on understanding rather than an unconscious inheritance from convention or momentum. Framework-free development gives you the foundation to make that choice deliberately.

Chapter 2: A Brief History of Software Abstraction

Understanding why we build software the way we do requires understanding how we got here. The modern landscape of framework-heavy development did not emerge by accident. It was shaped by real problems, technological shifts, and economic pressures that favored certain approaches over others. Tracing this history reveals patterns that help us evaluate current trends and anticipate future directions.

From Assembly to Libraries to Frameworks

Software development began without abstractions. Early programmers wrote directly in machine code, manipulating binary values to control hardware. Assembly languages provided symbolic names for opcodes and memory addresses, but the programmer still managed every register, every memory location, every instruction sequence. This was necessary because there was nothing else: no operating system to abstract away hardware details, no standard libraries to provide common functionality.

Operating systems introduced the first major layer of abstraction. System calls like read, write, open, and close hid the complexity of device drivers and disk controllers behind simple interfaces. Programmers could write portable code that worked across different hardware without knowing the specifics of each device. Standard libraries added functions for string manipulation, mathematical operations, memory allocation, and input/output formatting. These were pure libraries: the programmer called them when needed, maintaining full control over program flow.

The 1970s and 1980s saw the rise of high-level languages like C, Pascal, Ada, and later C++. These languages abstracted away low-level details through type systems, memory management utilities, and structured programming constructs. Object-oriented programming introduced classes, inheritance, and polymorphism as abstractions for organizing code around data and behavior.

Design patterns emerged as reusable solutions to common design problems. All of these were tools that programmers used voluntarily to structure their code better. They did not dictate program architecture or invert control flow.

The shift toward frameworks began in earnest with graphical user interface toolkits. Building a GUI application required managing windows, events, menus, and widgets across different operating systems and display technologies. Frameworks like Motif, Tkinter, and later Java's AWT/Swing provided event-driven architectures where the framework owned the main loop and called application code in response to user actions. This inversion of control made sense for GUI development because the application fundamentally reacts to external events. The framework's role as event dispatcher was natural and necessary.

Web development followed a similar trajectory. In the mid-1990s, building a web application meant writing CGI scripts or using server-side languages like PHP and ASP to generate HTML dynamically. Each request spawned a new process or thread that executed the script from start to finish. There was no framework, no structure, no conventions. Developers organized their code however they wanted, which led to inconsistent, hard-to-maintain applications as projects grew in complexity.

The Spring Revolution and Its Aftermath

The early 2000s marked a turning point with the rise of opinionated frameworks that imposed structure on entire application domains. On the Java side, Spring emerged as a comprehensive framework for enterprise application development. It provided dependency injection, transaction management, MVC web layers, data access abstractions, and integration with messaging systems. Spring Boot later automated much of the configuration through convention-based auto-configuration, allowing developers to create production-ready services with minimal setup.

Spring's success was not accidental. It solved real problems that Java developers faced: verbose boilerplate code for JDBC connections and transaction management, complex XML configuration for enterprise beans, inconsistent patterns across different parts of large applications. By providing a unified architecture and a rich ecosystem of modules, Spring reduced the cognitive burden of building enterprise systems. Teams could focus on business logic while the framework handled infrastructure concerns.

The same era saw similar movements in other ecosystems. Ruby on Rails, released in 2004, introduced convention over configuration as a guiding principle. Instead of requiring explicit configuration for every aspect of the application, Rails assumed sensible defaults based on naming conventions. Models corresponded to database tables with predictable names. Controllers handled HTTP requests for resources they were named after. Views rendered templates in standard directories. This convention-based approach dramatically accelerated development for standard web applications.

Django, released in 2005, took a similar batteries-included approach for Python. It provided an ORM, authentication system, admin interface, templating engine, and URL routing out of the box. The framework's philosophy was explicit: it should be easy to build complete web applications quickly without assembling disparate components from different projects.

These frameworks changed how developers thought about application architecture. Instead of designing their own structures and wiring their own components, developers learned to work within the framework's expectations. New programmers were taught framework conventions as fundamental programming knowledge. Job postings listed framework proficiency as a primary requirement. The frameworks became the de facto standard for building applications in their respective ecosystems.

This shift had clear benefits: faster development, consistent patterns across projects, shared tooling and expertise, mature ecosystems of plugins and extensions. But it also created dependencies that many teams did not fully appreciate until problems emerged. Frameworks accumulated features to serve broader audiences, making them larger and more complex. Internal APIs changed between versions, breaking applications that relied on implementation details. Debugging required understanding framework internals rather than just application code. Performance optimization meant working around framework limitations rather than designing systems optimized for specific workloads.

The Rise of Microframeworks and Minimalist Languages

By the 2010s, a counter-movement began as developers pushed back against framework bloat. Microframeworks emerged that provided minimal structure while leaving architectural decisions to the developer. Flask for Python offered

routing and request handling without imposing an ORM, templating engine, or project structure. Express for Node.js provided middleware composition and routing without dictating how applications should be organized. These microframeworks inverted control just enough to simplify web server boilerplate while leaving the rest of the architecture open.

At the same time, new programming languages emerged with philosophies that favored minimalism and explicit design. Go, released in 2009, included a comprehensive standard library that made external frameworks largely unnecessary for many use cases. The net/http package provided a fully functional HTTP server. The encoding/json package handled serialization. The database/sql package offered connection pooling and query execution. Developers could build production services using only the standard library and carefully chosen third-party libraries for specific needs.

Go's design philosophy emphasized simplicity, readability, and explicit behavior over clever abstractions. No inheritance hierarchies, no generics for many years, no macros, no operator overloading. The language forced developers to write straightforward code that was easy to understand and reason about. This philosophy extended to the ecosystem: popular Go projects tend to be libraries rather than frameworks, following the principle that the application should control its own structure.

Rust, stabilized in 2015, took a different approach but arrived at similar conclusions. Rust's ownership model and compile-time safety guarantees eliminate entire classes of bugs that runtime-heavy frameworks must handle defensively. This allows Rust applications to be leaner and faster than equivalent Java or Python services. While Rust does have web frameworks like Actix and Axum, many production Rust services use minimal abstractions, relying on the language's type system and standard library for structure rather than framework conventions.

The microservices architecture trend also encouraged simpler, more focused applications. Instead of one massive monolith built on a heavyweight framework, systems were decomposed into smaller services with single responsibilities. Each service could choose its own technology stack optimized for its specific workload. This environment favored lightweight libraries and minimal frameworks that added value without imposing unnecessary structure.

Current Trends: Back to Basics

As of the mid-2020s, the software development landscape shows clear signs of a pendulum swing back toward minimalism and explicit design. Several factors are driving this shift.

Performance pressure has intensified as systems scale to handle billions of requests per day. Framework overhead that was acceptable for small applications becomes unacceptable at massive scale. Cloudflare's migration from its NGINX-based FL1 proxy to Pingora, a Rust-built system, delivered seventy percent less CPU usage and sixty-seven percent less memory consumption while handling over one trillion HTTP requests per day [6]. This is not a marginal improvement; it represents hundreds of millions of dollars in infrastructure savings and measurable latency reduction for end users.

Security concerns have grown as software supply chain attacks become more common and damaging. The Log4j vulnerability demonstrated how deeply nested dependencies can expose organizations to risk. Framework-free systems with minimal, well-audited dependencies have a smaller attack surface and are easier to verify for security issues. Teams are becoming more selective about what they include in their dependency trees.

Developer experience has evolved as tooling improves. Modern IDEs provide excellent code completion, refactoring, and navigation support that reduces the need for frameworks to enforce structure through conventions. Linters, formatters, and static analysis tools catch errors and enforce style guidelines without requiring framework-specific patterns. Developers can write clean, consistent code using explicit design rather than relying on framework auto-configuration.

Operational maturity has reduced the value of framework-provided operational features. Container orchestration platforms like Kubernetes handle service discovery, load balancing, health checking, and configuration management at the infrastructure level. Observability platforms provide distributed tracing, metrics collection, and log aggregation that work across any application regardless of framework. Features that frameworks once bundled as selling points are now available as platform capabilities.

The result is a growing appreciation for systems that are simple, transparent, and under developer control. Frameworks have not disappeared; they remain valuable tools for specific use cases. But the default assumption is

shifting from "what framework should we use" to "what is the minimal set of abstractions needed to solve this problem." This mindset aligns directly with the principles explored throughout this book.

Understanding this history helps you evaluate frameworks more critically. When you encounter a new framework promising to solve your problems, you can ask whether it addresses a genuine need or simply adds another layer of indirection. You can recognize when framework conventions serve your application and when they constrain it. And you can appreciate that the skills to build without frameworks are not relics of an earlier era; they are increasingly valuable in a landscape that rewards simplicity, performance, and control.

The next chapters will show you how to apply this understanding in practice, starting with the architectural principles that make framework-free development viable and robust.

Chapter 3: Core Architectural Principles

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/framework-freesoftwareengineering.

SOLID Principles Without the Buzzword Theater

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/framework-freesoftwareengineering.

Clean Architecture in Practice

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/framework-freesoftwareengineering.

Hexagonal Architecture and Port/Adapter Patterns

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/framework-freesoftwareengineering.

The Dependency Rule and Inversion

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/framework-freesoftwareengineering.

Chapter 4: Modular Design and Separation of Concerns

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/framework-freesoftwareengineering.

Defining Module Boundaries That Hold

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/framework-freesoftwareengineering.

Information Hiding and Encapsulation

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/framework-freesoftwareengineering.

Package Design Across Languages

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/framework-freesoftwareengineering.

Feature Modules vs Layered Modules

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/framework-freesoftwareengineering.

Chapter 5: Dependency Management Without Frameworks

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/framework-freesoftwareengineering.

Constructor Injection by Hand

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/framework-freesoftwareengineering.

Service Locators and When to Avoid Them

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/framework-freesoftwareengineering.

Factory Patterns for Flexible Composition

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/framework-freesoftwareengineering.

Wiring Up a Real Application

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/framework-freesoftwareengineering.

Chapter 6: Configuration, Environment, and Bootstrapping

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/framework-freesoftwareengineering.

The Twelve-Factor Approach to Configuration

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/framework-freesoftwareengineering.

Environment Variables, Files, and Hierarchies

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/framework-freesoftwareengineering.

Secrets Management Without Ceremony

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/framework-freesoftwareengineering.

Bootstrapping: From Zero to Ready

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/framework-freesoftwareengineering.

Chapter 7: API Design and Communication Patterns

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/framework-freesoftwareengineering.

RESTful APIs Built From Scratch

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/framework-freesoftwareengineering.

HTTP Semantics Done Right

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/framework-freesoftwareengineering.

Protocol Buffers and gRPC Without Magic

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/framework-freesoftwareengineering.

Event-Driven Communication Patterns

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/framework-freesoftwareengineering.

Chapter 8: Error Handling, Logging, and Observability

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/framework-freesoftwareengineering.

Error Types, Propagation, and Recovery

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/framework-freesoftwareengineering.

Structured Logging as a First-Class Concern

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/framework-freesoftwareengineering.

Metrics, Tracing, and Distributed Observability

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/framework-freesoftwareengineering.

Health Checks and Readiness Probes

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/framework-freesoftwareengineering.

Chapter 9: Persistence and Data Access Without ORMs

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/framework-freesoftwareengineering.

Why ORMs Often Fail You

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/framework-freesoftwareengineering.

Writing SQL That Scales

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/framework-freesoftwareengineering.

The Repository Pattern Implemented

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/framework-freesoftwareengineering.

Migrations, Connection Pooling, and Transactions

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/framework-freesoftwareengineering.

Chapter 10: Concurrency, Async Programming, and Performance

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/framework-freesoftwareengineering.

Concurrency Models Across Languages

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/framework-freesoftwareengineering.

Async/Await Without Framework Magic

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/framework-freesoftwareengineering.

Locks, Channels, and Lock-Free Patterns

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/framework-freesoftwareengineering.

Profiling and Performance Tuning

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/framework-freesoftwareengineering.

Chapter 11: Security in Minimalist Architectures

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/framework-freesoftwareengineering.

Authentication Without Framework Helpers

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/framework-freesoftwareengineering.

Authorization Models From Scratch

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/framework-freesoftwareengineering.

Encryption, Hashing, and Key Management

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/framework-freesoftwareengineering.

Input Validation and Attack Surface Reduction

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/framework-freesoftwareengineering.

Chapter 12: Testing Strategies for Framework-Free Systems

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/framework-freesoftwareengineering.

Why Framework-Free Code Tests Better

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/framework-freesoftwareengineering.

Unit Testing With Hand-Wired Dependencies

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/framework-freesoftwareengineering.

Integration Tests That Are Fast and Reliable

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/framework-freesoftwareengineering.

Property-Based Testing and Fuzzing

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/framework-freesoftwareengineering.

Chapter 13: Build, Deploy, and Operate

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/framework-freesoftwareengineering.

Build Tools and Dependency Resolution

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/framework-freesoftwareengineering.

Container Images That Are Small and Safe

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/framework-freesoftwareengineering.

Deployment Strategies and Rollbacks

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/framework-freesoftwareengineering.

Scaling Patterns and Operational Runbooks

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/framework-freesoftwareengineering.

Chapter 14: Case Studies: Real-World Framework-Free Systems

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/framework-freesoftwareengineering.

A High-Throughput Go API Service

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/framework-freesoftwareengineering.

A Rust CLI Tool at Scale

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/framework-freesoftwareengineering.

A Minimalist Node.js Microservice

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/framework-freesoftwareengineering.

Lessons From Systems That Refused the Framework

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/framework-freesoftwareengineering.

Chapter 15: Conclusion: Choosing Your Path

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/framework-freesoftwareengineering.

The Decision Matrix: Framework vs Framework-Free

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/framework-freesoftwareengineering.

Hybrid Approaches That Actually Work

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/framework-freesoftwareengineering.

Signs Your Framework Has Become a Prison

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/framework-freesoftwareengineering.

Final Thoughts on Engineering Judgment

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/framework-freesoftwareengineering.

References

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/framework-freesoftwareengineering.