MODERN CODE REVIEW

IN THE AI ERA

A Comprehensive Guide to Building Better Software Through Collaborative Review and Intelligent Automation

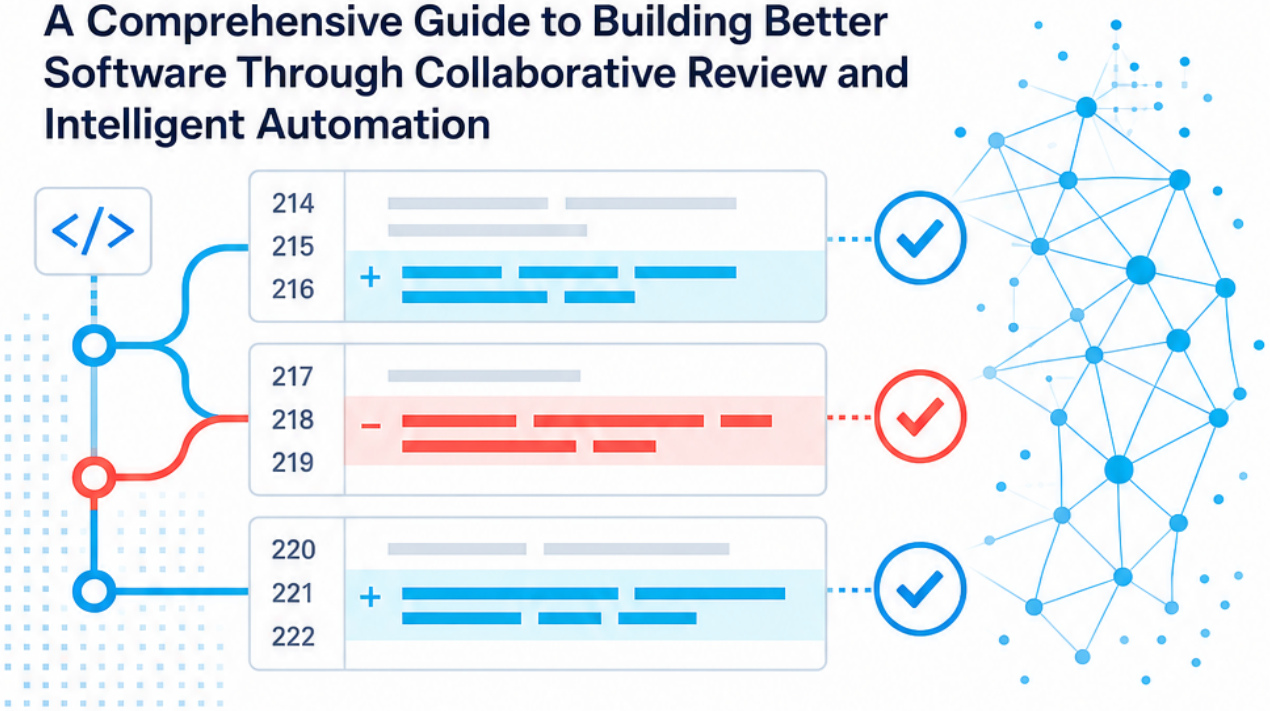

CAMILLE BERNARD

Modern Code Review in the AI Era

A Comprehensive Guide to Building Better Software Through Collaborative Review and Intelligent Automation

Steve Publications

This book is available at https://leanpub.com/moderncodereviewintheaiera

This version was published on 2026-07-28

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.

Contents

A Comprehensive Guide to Building Better Software Through Collaborative Review and Intelligent Automation

This book is for software engineers who want to review code with precision, confidence, and judgment. It traces the evolution of code review from formal inspections to modern pull requests, explains what makes software maintainable and secure, and provides systematic guidance for reviewing everything from unit tests to distributed systems architectures. The second half addresses the arrival of AI coding assistants: how they transform review workflows, what new risks they introduce, how to detect hallucinations and subtle bugs in AI-generated code, and how organizations can govern responsible AI-assisted development. Every chapter includes production-quality examples, real-world scenarios, decision frameworks, and industry best practices drawn from open-source projects and leading technology companies.

Introduction: The Review Imperative

A Bug That Cost Billions: Why Review Matters

On July 19, 2024, a single faulty configuration update from cybersecurity company CrowdStrike crashed approximately 8.5 million Windows systems worldwide. Airlines grounded flights. Hospitals diverted patients. Banks suspended transactions. News organizations went dark. The global economic impact was estimated at billions of dollars in lost productivity alone, making it one of the largest IT outages in history.

The root cause was not a sophisticated cyberattack or a complex distributed system failure. It was a logic error in a configuration file for CrowdStrike's Falcon Sensor security software. A new update included a rule that referenced an array index beyond its bounds. When the sensor attempted to evaluate this rule on every Windows machine, it triggered a blue screen crash. The faulty update had passed through CrowdStrike's internal testing pipeline, but the test environment apparently did not cover the specific code path that caused the crash in production. A more thorough review of the configuration logic, combined with better canary deployment practices, could have contained the damage to a small subset of systems rather than allowing it to cascade globally.

This incident demonstrates a fundamental truth about software: even organizations with vast resources, sophisticated tooling, and experienced engineers are vulnerable to simple mistakes that slip through their defenses. Code review exists for exactly this reason: to give every change a second pair of eyes before it reaches production.

The economics are stark. Industry research has long documented that defects found in production cost orders of magnitude more than those caught during development. Conservative estimates place the multiplier between thirty and one hundred times, depending on severity and scope. A bug discovered in requirements might cost a few hundred dollars to fix. The same bug, once deployed and affecting customers, can trigger incident response, customer churn, regulatory scrutiny, and reputational damage that runs into

millions. Code review is not bureaucracy; it is one of the highest-return investments an engineering organization can make.

Multiple empirical studies confirm this. Research on defect detection across development phases shows unit testing catches approximately 25 percent of defects, functional testing around 35 percent, and integration testing about 45 percent. In contrast, design and code inspections detect between 55 and 60 percent of defects, including many that automated tests never encounter because they require understanding developer intent, architectural coherence, and security implications. One landmark study at AT&T involving more than two hundred engineers reported a 14 percent productivity increase and a 90 percent reduction in defects after introducing systematic code review.

What Code Review Actually Is (And Is Not)

Code review is the practice of having non-authors examine proposed changes to source code before they are integrated into the main codebase. At its best, it serves multiple purposes simultaneously:

- It catches bugs, security vulnerabilities, and design flaws early, when they are cheap to fix.
- It spreads knowledge across the team, preventing bus factors and silos of expertise.
- It enforces consistency in style, architecture, and conventions without requiring a dictator.
- It provides a forum for discussing trade-offs and making architectural decisions collaboratively.
- It accelerates onboarding, as junior engineers learn by reading good code and receiving feedback.

What code review is not:

- It is not a gatekeeping exercise designed to assert authority or demonstrate superiority.
- It is not a substitute for automated testing, static analysis, or CI pipelines, which should catch mechanical errors before humans ever see the diff.
- It is not a design document, though it can surface design issues that warrant deeper discussion.

- It is not an opportunity to nitpick whitespace, naming preferences, or stylistic choices that do not affect correctness, readability, or maintainability.

The distinction matters because treating review as any of these non-things corrupts the culture around it. When engineers fear review, they write defensively rather than creatively, hide problems rather than expose them, and game metrics rather than improve quality. The goal is to build a culture where every engineer views review as collaboration, not judgment.

The AI Inflection Point

Now consider what happens when you introduce artificial intelligence into this ecosystem.

By early 2025, GitHub Copilot had more than twenty million users across over fifty thousand organizations, including ninety percent of the Fortune 100. Microsoft reported that between 20 and 30 percent of code committed in its own repositories was AI-generated. Google published similar figures. A study published in Science in January 2026 found that in the United States, the share of new code relying on AI assistance rose from 5 percent in 2022 to 29 percent by early 2025, compared with just 12 percent in China.

This is not a trend. It is a transformation.

The implications for code review are profound and dual-edged. On the positive side, AI tools can accelerate the mechanical parts of review: checking style, flagging obvious bugs, suggesting tests, and summarizing large diffs. GitHub's own Copilot Code Review feature, which reached general availability in April 2025, had completed more than sixty million reviews by March 2026 and now covers over 20 percent of all pull requests on the platform. Teams report shipping code faster and spending less time on repetitive review comments.

On the negative side, AI introduces risks that did not exist before. A comprehensive study published in October 2025 analyzed 7,703 files explicitly attributed to AI tools across public GitHub repositories and found that Python code generated by these models exhibited vulnerability rates between 16 and 18 percent, with SQL injection, OS command injection, and code injection among the most common issues. Veracode's 2025 GenAI Code Security Report, which tested over one hundred large language models, found that only 55 percent of

AI-generated code passed security checks. Java was particularly problematic, with a security pass rate of just 29 percent. Models failed against cross-site scripting in 86 percent of test cases and log injection in 88 percent.

Then there are hallucinations, where models confidently recommend packages that do not exist, APIs that were never released, or functions with incorrect signatures. A USENIX Security 2025 study on package hallucinations tested sixteen models across approximately nineteen thousand prompts and found an average hallucination rate of 19.6 percent, with open-source models averaging 21 percent and commercial models averaging 5 percent. The researchers generated more than two hundred thousand unique fictitious package names and warned that developers who blindly trust these suggestions risk installing malicious software through dependency confusion attacks.

The central thesis of this book is that AI has not made code review obsolete. It has made it more important than ever, while fundamentally changing what reviewers must look for and how they approach their work. Human judgment remains essential because:

- AI tools do not understand business context, domain constraints, or architectural intent the way experienced engineers do.
- AI-generated code can be functionally correct while being architecturally wrong, inefficient, insecure, or unmaintainable.
- AI models produce plausible-sounding output that masks errors, making superficial review more dangerous, not less.
- The responsibility for what ships to production ultimately rests with humans, not models.

The organizations that thrive in the AI era will be those that combine intelligent automation with thoughtful human oversight, using AI as a powerful assistant while maintaining rigorous standards for quality, security, and maintainability.

How to Use This Book

This book is organized to build knowledge progressively, from foundations through advanced practices and AI-era transformations. Here is how the chapters fit together:

The Introduction you are reading now establishes why code review matters and frames the challenges and opportunities introduced by AI.

Chapter 1, "Origins and Evolution of Code Review," traces the history from formal inspections in the 1970s through modern pull request workflows, showing how past lessons inform current practice.

Chapter 2, "Principles of Maintainable, Secure, High-Quality Software," establishes the quality criteria that all review should uphold: readability, complexity management, correctness, security, performance, and observability. This is a reference chapter you will return to.

Chapter 3, "Psychology, Culture, and Collaborative Review," addresses the human side of code review, covering psychological safety, constructive feedback, reviewer fatigue, and building norms that improve both code and team dynamics.

Chapter 4, "Git-Based Workflows and Pull Request Mechanics," provides a thorough reference on branching strategies, PR lifecycle management, commit hygiene, and tooling choices for GitHub, GitLab, and Bitbucket.

Chapter 5, "Review Automation and Static Analysis," covers linters, formatters, static analyzers, type checkers, secret scanners, and how to configure them as an effective review safety net.

Chapter 6, "Testing Strategies Reviewed," teaches you how to evaluate test suites accompanying changes, including unit tests, integration tests, property-based testing, and mutation testing.

Chapter 7, "Security Review," equips you with systematic approaches for spotting vulnerabilities: injection flaws, authentication bugs, cryptographic misuse, supply chain risks, and cloud misconfigurations.

Chapter 8, "Performance, Scalability, and Reliability Review," covers reviewing code for performance characteristics, concurrency safety, resource management, database access patterns, and fault tolerance.

Chapter 9, "Architectural Review," addresses evaluation of architectural decisions: module boundaries, interface design, domain modeling, and system-level trade-offs.

Chapter 10, "Domain-Specific Review Practices," provides specialized guidance for distributed systems, cloud-native applications, APIs and microservices, frontend applications, mobile development, and infrastructure as code.

Chapter 11, "Documentation, Observability, and Operational Readiness," covers review of non-code artifacts: documentation quality, logging strategy, metrics, tracing, runbooks, and deployment readiness.

Chapter 12, "AI-Assisted Code Review: The New Paradigm," provides comprehensive treatment of how AI and large language models are transforming code review, including reviewing AI-generated code, detecting hallucinations, using AI as a review copilot, prompt engineering for reviewers, and integrating AI into CI/CD pipelines.

Chapter 13, "Governance, Policy, and Organizational Strategy," addresses the organizational layer: establishing review policies, metrics that matter and ones to avoid, training programs, AI governance frameworks, and scaling practices across large organizations.

The Conclusion synthesizes the book's themes and provides a practical checklist for modern review excellence.

You can read this book cover to cover, or use it as a reference, jumping to chapters relevant to your current needs. Each chapter is self-contained but builds on earlier material, so foundational concepts are introduced before they are reused. Code examples throughout are production-quality and fully compilable, using modern versions of Python, Go, Java, C#, Rust, TypeScript, JavaScript, Kotlin, and C++. Real-world pull requests, review comments, CI/CD configurations, and case studies illustrate concepts in context.

Let us begin.

Chapter 1: Origins and Evolution of Code Review

Before Computers: Engineering Inspection Heritage

The practice of systematic inspection predates software engineering by centuries. In the late nineteenth century, mechanical and civil engineers developed rigorous review processes for bridge designs, engine blueprints, and architectural plans. A design would be examined by multiple specialists, each checking for failures within their domain: structural integrity, material stress tolerances, thermal expansion effects. Errors caught on paper were cheap to fix; errors discovered after construction could be catastrophic.

This heritage directly influenced early software engineering. The first practitioners of programming recognized that writing code without systematic review was analogous to building bridges without structural analysis. As software systems grew in complexity and criticality, the need for disciplined inspection became unavoidable.

The IBM/Boeing Era: Formal Inspections Born

The modern history of code review begins with Michael Fagan at IBM's Thomas J. Watson Research Center in the early 1970s. Fagan was tasked with improving the quality of software developed for mainframe systems, where defects were expensive and often discovered only after deployment. His experiments produced a formal process that would become known as the "Fagan inspection."

Fagan's approach treated code review not as an informal peer discussion but as a structured, multi-phase process with defined roles, checklists, and metrics. A typical Fagan inspection involved:

1. **Planning**: Selecting the appropriate reviewers and scheduling the inspection.

2. **Overview meeting**: The author briefly explained the code's purpose and design to give reviewers context.
3. **Individual preparation**: Each reviewer examined the code independently, marking defects on paper copies.
4. **Inspection meeting**: Reviewers presented their findings in a structured discussion, with defects recorded by a designated scribe.
5. **Rework**: The author fixed identified defects.
6. **Follow-up**: A second inspection might occur if many defects were found, confirming fixes without re-examining everything.

Fagan's experiments at IBM showed dramatic improvements in defect detection and productivity. His seminal 1976 paper, "Design and Code Inspections to Reduce Errors in Program Development," documented defect removal rates of 50 to 90 percent during the inspection phase itself, far exceeding what testing alone achieved.

Around the same period, other researchers and practitioners were exploring similar ideas. Glenford Myers and Tom DeMarco published influential work on software quality assurance. Boeing adopted formal review processes for flight-critical software, recognizing that aviation demanded higher standards than typical commercial applications. The Defense Advanced Research Projects Agency (DARPA) funded research into software reliability, further legitimizing inspection as a core engineering practice.

Fagan Inspections and the Software Quality Movement

By the 1980s, Fagan inspections had spread beyond IBM to major technology companies and government contractors. They became a standard component of high-assurance software development, particularly in aerospace, defense, and medical device industries where failures could cost lives.

The inspection movement also influenced standards bodies. The IEEE published guidelines for software inspections. ISO incorporated review requirements into its quality management standards. CMM (Capability Maturity Model), developed at the Software Engineering Institute at Carnegie Mellon University, treated formal reviews as a key indicator of organizational maturity.

However, Fagan inspections had significant limitations that would eventually drive evolution toward lighter-weight approaches:

- They were time-intensive, requiring hours of preparation and meeting time for each review session.
- The formal structure could feel rigid and bureaucratic to engineers accustomed to more informal collaboration.
- They worked best for relatively small changes; inspecting large modules became unwieldy.
- Printed code copies, standard at the time, were impractical for rapidly changing projects.

These limitations did not invalidate inspections as a concept. Rather, they highlighted a tension that persists today: how to maintain rigor without sacrificing agility.

From Waterfall to Agile: Lightweight Review Emerges

The 1990s brought two developments that reshaped software development and, with it, code review. The first was the rise of version control systems beyond simple file locking. CVS (Concurrent Versions System) and later Subversion introduced diff-based change tracking, enabling reviewers to see exactly what had changed rather than re-examining entire files.

The second was the emergence of agile methodologies. The Agile Manifesto, published in 2001, emphasized "individuals and interactions over processes and tools" and "working software over comprehensive documentation." These values challenged heavyweight inspection processes that felt at odds with rapid iteration.

Agile teams developed lighter-weight review practices adapted to shorter cycles:

- **Walkthroughs**: The author led reviewers through the code in real time, explaining decisions and answering questions. Less formal than inspections, walkthroughs were faster and more conversational.
- **Over-the-shoulder reviews**: Two engineers sat together, one sharing their screen while the other reviewed changes as they were written. This was essentially continuous review embedded into pair programming.
- **Email-based review**: Some teams circulated diffs via email, with reviewers adding comments inline. This enabled asynchronous review without requiring a formal meeting.

These approaches preserved the core benefit of multiple eyes on code while reducing ceremony. They also reflected a cultural shift: review was no longer primarily about defect detection in high-assurance contexts but about collaborative improvement in fast-moving development environments.

Version Control as Collaboration: CVS, SVN, and the Birth of Diff-Based Review

Version control systems evolved from tools for tracking changes to platforms for collaboration. In the early days of CVS and Subversion, review was an adjunct activity: developers committed code and then circulated diffs via email or issue trackers. Reviewers read patches, added comments, and requested changes. The author applied fixes in subsequent commits.

This workflow had flaws. Review happened after commit, meaning defective code could sit in the repository until someone noticed and fixed it. There was no clear mechanism for tracking review status: who had reviewed, what they thought, whether approval was required before merging.

Some organizations built custom solutions. Google developed internal tools for code review as its engineering organization grew. The Linux kernel community used mailing lists to discuss proposed patches, with Linus Torvalds and subsystem maintainers reviewing submissions before integration. This model, while effective for a distributed open-source project, did not scale well to commercial development.

The critical insight was that version control and review should be integrated. A change should not be merged until it had been reviewed and approved. The tooling should make this explicit, tracking who reviewed what and enforcing organizational policies.

Git and the Pull Request Revolution

Git, created by Linus Torvalds in 2005 for Linux kernel development, eventually enabled a new paradigm through platforms like GitHub, launched in 2008. The pull request (PR), as GitHub called it, or merge request (as GitLab termed it), unified version control, review, and collaboration into a single workflow:

1. A developer creates a branch from the main codebase.

2. They make changes and push commits to that branch.
3. They open a pull request, which displays the diff, links to related issues, and invites reviewers.
4. Reviewers examine the changes, leave inline comments, and approve or request modifications.
5. The author addresses feedback with additional commits.
6. Once approved (and typically after automated checks pass), the PR is merged into the main branch.

This model solved many problems of earlier workflows:

- Review happened before merge, preventing unreviewed code from entering the main branch.
- Discussion was threaded and visible to the entire team, creating a record of decisions.
- Status was explicit: open, in review, approved, merged.
- Automation could be integrated: CI pipelines ran tests on every push, blocking merge until checks passed.

The pull request became the dominant collaboration mechanism for software development. GitHub's rise to becoming the central hub for both open-source and enterprise development cemented this model. By the mid-2010s, most professional engineering teams were using PR-based review, whether on GitHub, GitLab, Bitbucket, or internal equivalents.

Modern code review tools added capabilities that Fagan could scarcely have imagined:

- **Inline commenting**: Reviewers annotate specific lines rather than writing separate reports.
- **Suggested changes**: Reviewers propose edits directly in the diff, which authors can accept with one click.
- **Automated checks**: CI pipelines run tests, linters, and security scans on every change.
- **Review assignments**: Tools route PRs to appropriate reviewers based on file ownership or expertise.
- **Metrics and analytics**: Organizations track review times, comment volumes, and defect rates to improve processes.

Lessons from History That Still Apply Today

Understanding the evolution of code review is not an academic exercise. The history reveals principles that remain valid regardless of tooling:

Multiple eyes are better than one. This is the fundamental insight behind every review methodology, from Fagan inspections to modern pull requests. Different reviewers notice different things, and no individual catches everything.

Early detection saves money. Defects found during review cost far less to fix than those discovered in testing or production. This economic reality has not changed.

Process should serve people, not the reverse. Fagan inspections were rigorous but cumbersome, leading teams to develop lighter alternatives. The best practices are those engineers actually follow, not those that look good on a maturity model.

Tooling amplifies culture. Good tools make review easier and more effective. Poor tools create friction that engineers work around. The pull request succeeded because it fit naturally into how developers already worked with version control.

Review is about more than bugs. Modern code review, like its predecessors, serves multiple purposes: knowledge sharing, quality enforcement, architectural alignment, and team learning. Treating it as purely a defect-detection mechanism underutilizes its value.

As we move into the AI era, these principles remain anchors. AI tools can accelerate mechanical aspects of review, but they cannot replace the human judgment, contextual understanding, and collaborative dynamics that make code review valuable. The next chapters explore how to apply these lessons in practice.

Chapter 2: Principles of Maintainable, Secure, High-Quality Software

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

Readability as a First-Class Requirement

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

The Complexity Tax: Cyclomatic, Cognitive, and Beyond

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

Correctness Through Defensiveness and Explicit Contracts

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

Security by Design: Threat Modeling in Review

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

Performance as a Continuous Concern

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

Observability: Code That Can Be Debugged in Production

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

The Maintainability Equation: Coupling, Cohesion, and Change Cost

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

Chapter 3: Psychology, Culture, and Collaborative Review

Ego Is the Enemy of Good Code

Building Psychological Safety in Review

The Art of the Constructive Comment

Reviewer Fatigue and Its Remedies

Size Matters: Why Small PRs Get Better Reviews

Norms, Charters, and Written Expectations

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

Conflict Resolution When Reviewers Disagree

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

Chapter 4: Git-Based Workflows and Pull Request Mechanics

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

Branching Strategies Compared: Git Flow, Trunk-Based, GitHub Flow

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

The Anatomy of a High-Quality Pull Request

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

Commit Hygiene and Atomic Changes

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

Reviewer Assignment and Load Balancing

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

Approvals, Rejections, and the Approval Quorum

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

Merge Strategies: Squash, Rebase, Merge Commits

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

GitHub Actions, GitLab CI, and Bitbucket Pipelines for Review Gating

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

Chapter 5: Review Automation and Static Analysis

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

The Automation Pyramid: What to Automate and When

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

Linting and Formatting: Enforcing Style Without Debate

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

Static Analysis Engines: SonarQube, CodeQL, Semgrep, and Others

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

Type Systems as Review Assistants

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

Secret Detection and Dependency Scanning

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

Building a CI Pipeline That Catches Problems Before Humans See Them

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

False Positives, Alert Fatigue, and Tool Tuning

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

Chapter 6: Testing Strategies Reviewed

What Makes a Test Worth Reviewing

Unit Tests: Isolation, Fakes, and Meaningful Assertions

Integration Tests: Boundaries, Contracts, and Real Dependencies

End-to-End Tests: When They Earn Their Cost

Property-Based Testing and Mutation Testing in Review

Performance and Load Tests as First-Class Code

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

The Test Smells Every Reviewer Should Know

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

Chapter 7: Security Review: Finding Vulnerabilities Before Attackers Do

Thinking Like an Attacker During Review

OWASP Top 10 Through the Reviewer's Eyes

Authentication, Authorization, and Session Management

Cryptography: What Not to Roll Yourself

Input Validation, Output Encoding, and Sanitization

API Security: Rate Limiting, Headers, and CORS

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

Supply Chain Security and Dependency Review

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

Secrets Management and Configuration Hardening

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

Simulated PR Review: Security-Focused

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

PR Description

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

Round 1 Review

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

Author Response to Round 1

Round 2 Review (After Fixes)

Author Response to Round 2

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

Final Review

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

Key Lessons from This Review

Chapter 8: Performance, Scalability, and Reliability Review

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

Complexity Analysis as a Review Skill

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

Memory Management and Leak Detection

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

Concurrency: Race Conditions, Deadlocks, and Lock-Free Design

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

Database Access Patterns and Query Performance

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

Caching Strategies and Their Pitfalls

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

Timeout, Retry, and Circuit Breaker Patterns

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

Load Testing Evidence in Pull Requests

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

Chapter 9: Architectural Review: Evaluating Structure at Scale

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

When Code Review Becomes Architecture Review

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

Module Boundaries and Dependency Direction

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

Interface Design: Stability, Clarity, and Backward Compatibility

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

Domain Modeling and Business Logic Placement

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

Event-Driven Architectures and Message Contracts

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

Migration Strategies and Incremental Refactoring

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

Technical Debt: Recognizing, Tracking, and Paying It Down

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

Simulated PR Review: Architectural Decision

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

PR Description

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

Round 1 Review

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

Architecture Discussion (Scheduled Meeting)

Round 2 Review (After Discussion and Updates)

Author Response and Final Notes

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

Key Lessons from This Review

Chapter 10: Domain-Specific Review Practices

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

Distributed Systems: Consistency, Partition Tolerance, and Timeouts

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

Cloud-Native Applications: Containers, Orchestration, and Config

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

API and Microservice Review: Contracts, Versioning, and Idempotency

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

Frontend Application Review: State Management, Accessibility, and Performance

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

Mobile Development: Platform Conventions, Lifecycle, and Offline Behavior

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

Infrastructure as Code: Terraform, Kubernetes Manifests, and Policy

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

Chapter 11: Documentation, Observability, and Operational Readiness

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

Documentation as Code: READMEs, API Docs, and Architecture Decision Records

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

Architecture

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

API Documentation

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

Contributing

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

Logging Strategy: Levels, Structure, and Signal-to-Noise

Metrics That Matter: RED, USE, and Business Signals

Distributed Tracing and Correlation IDs

Runbooks and Incident Response Readiness

Feature Flags, Rollout Strategy, and Canary Analysis

Chapter 12: AI-Assisted Code Review: The New Paradigm

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

How AI Coding Assistants Work (And Don't Work)

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

The Unique Risks of Reviewing AI-Generated Code

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

Detecting Hallucinations: Fake APIs, Wrong Semantics, Confident Lies

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

Using AI as a Review Copilot: Tools, Prompts, and Workflows

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

Validating AI Suggestions: Correctness, Security, and Maintainability

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

Case Studies: Humans Catching AI Errors

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

Case Study 1: Hallucinated API Usage That Compiled But Failed at Runtime

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

Case Study 2: Subtle Race Condition in AI-Generated Concurrent Code

Case Study 3: AI-Generated SQL with Silent Logic Error

Case Study 4: Security Vulnerability in AI-Generated Authentication Code

Prompt Engineering for Better AI Review Output

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

Integrating AI into CI/CD Pipelines

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

Measuring AI Impact: What Changes When Machines Help Review?

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

Chapter 13: Governance, Policy, and Organizational Strategy

Writing a Code Review Policy That Engineers Actually Follow

Metrics That Drive Improvement (And Ones That Backfire)

Training Programs and Onboarding New Reviewers

AI Governance: Acceptable Use, Data Privacy, and IP Concerns

Scaling Review Across Hundreds of Teams

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

When to Require Human Signoff vs. Automated Approval

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

The Future Trajectory: Where Code Review Is Headed

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

Conclusion: Judgment in the Age of Intelligence

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

The Enduring Value of Human Review

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

Combining Automation, AI, and Judgment into a Coherent Practice

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

A Checklist for Modern Review Excellence

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

Final Thoughts on Building Software That Lasts

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

References

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