

OpenSpec Deep Dive: Spec-Driven Development Architecture & Practice in AI-Assisted Programming

Created on January 11, 2026 Source: <https://github.com/Fission-AI/OpenSpec>

Tags: #AI Development #OpenSpec #Spec-Driven Development #SDD #AI Agents

Introduction: The Paradigm Shift in AI-Assisted Programming

As Large Language Model (LLM) capabilities grow exponentially, the software engineering field is experiencing profound transformation. Developers have shifted from traditional character-by-character coding to intent-driven orchestration. However, this shift hasn't come without costs. Current AI programming practices are plagued by a phenomenon called "vibe coding"—developers interact with AI through unstructured natural language conversations, with requirements scattered across lengthy chat logs, lacking persistence and systematization¹.

The Predicament of "Vibe Coding"

While "vibe coding" lowers the barrier to entry, its unsustainability in complex projects becomes increasingly apparent. When context windows fill up, AI often exhibits "amnesia" symptoms, leading to logic gaps, code regression, and severe hallucinations. Research shows that when context usage exceeds 40%, AI performance significantly degrades, with previous requirement details forgotten or tampered with in subsequent conversations². Moreover, this pattern makes code reviews extremely difficult, as reviewers cannot compare generated code against clear, written requirement specifications.

OpenSpec's Genesis & Mission

Against this backdrop, OpenSpec emerged as a standardized "Spec-Driven Development (SDD)" framework. It's not merely a tool but an engineering methodology aiming to solve AI programming's context loss and uncontrollability through the principle of "structure before code"³.

OpenSpec's core value proposition lies in its "brownfield-first" strategy. Unlike many AI tools focused on building new applications from scratch (greenfield, 0→1), OpenSpec is specifically designed to handle existing codebase evolution (1→n). By physically isolating current system state (Source of Truth) from proposed changes (Change Proposals) in the filesystem, it ensures change atomicity and auditability⁴. This architecture transforms AI agents from black boxes generating code into intelligent collaborators capable of understanding, planning, and executing explicit tasks.

Core Architecture & Design Philosophy

OpenSpec's architecture embodies the fusion of minimalism and pragmatism. It rejects complex database dependencies, instead adopting a Markdown-based filesystem storage solution. This design ensures specifications can reside alongside source code in version control systems (like Git), becoming "living documentation"¹.



structure serving as the AI agent's "external long-term memory." This structure is not only human-readable but optimized for machine parsing⁴.

Root Directory Structure Overview

A typical OpenSpec project contains these core components:

```
project_root/
├── openspec/
│   ├── AGENTS.md          # AI agent instruction sets &
behavioral guidelines
│   └── project.md         # Global project context, tech stack &
specs
│   ├── specs/            # Current system's "source of truth"
│   │   ├── auth-login/  # Capability-organized spec folders
│   │   │   └── spec.md
│   │   └── payment/
│   │       └── spec.md
│   └── changes/         # Active change proposals (workspace)
│       └── add-oauth-login/
│           ├── proposal.md # Change intent & high-level design
│           ├── design.md  # Technical architecture decisions
│           ├── tasks.md   # Atomic implementation task checklist
│           └── specs/     # Spec deltas (added/modified/removed)
```

Source of Truth (specs/)

The specs/ directory is OpenSpec's core asset repository. Organized by functional modules (capabilities), it stores the system's current business logic and technical constraints. For example, auth-login/spec.md might detail login flow input validation rules, error handling mechanisms, and security requirements. When AI agents receive modification tasks, they first retrieve from this directory to understand the system's existing behavior. This mechanism effectively prevents AI from accidentally breaking one feature's logic when modifying another, solving the "catastrophic forgetting" problem in long conversations².

Change Workspace (changes/)

OpenSpec introduces a change management model similar to Git branches. All feature requests or bug fixes initially manifest as independent subfolders under changes/. This isolation is crucial—it allows developers and AI to iteratively refine requirements without polluting the main spec repository. Only after changes are implemented and accepted are these temporary spec deltas merged back into the main specs/ directory⁴.

Agent Interface (AGENTS.md)

The AGENTS.md file is aptly called the "README for Robots." Unlike human-facing README.md, it contains fine-grained instructions for AI models, guiding how to read project context, format output, and follow OpenSpec's workflow state machine⁵. This file's existence gives OpenSpec high universality—even AI tools without native OpenSpec integration can follow specs by reading this file⁶.

State Machine Model & Lifecycle

OpenSpec enforces a strict software evolution state machine, dividing the development process into four irreversible phases:

1. Proposal	/openspec:proposal	proposal.md, tasks.md, design.md	Clarify intent, formulate plans, decompose tasks ³
2. Definition	openspec validate	specs/ (Deltas)	Write specific spec deltas (added, modified, removed) ⁶
3. Apply	/openspec:apply	Source code changes	AI writes code based on task checklist and specs ⁷
4. Archive	/openspec:archive	Merged specs/	Solidify changes into permanent docs, clean workspace ⁴

This structured lifecycle ensures documentation never lags behind code, fundamentally solving the “documentation rot” chronic ailment in software engineering⁷.

Environment Configuration & System Initialization

As a Node.js-based command-line tool (CLI), OpenSpec’s installation and configuration are designed to be as lightweight as possible, requiring no complex dependency chains or SaaS platform API keys—crucial for privacy-conscious enterprise environments³.

Installation Prerequisites & Steps

To run OpenSpec, the host system must have Node.js installed, with version recommended at v20.19.0 or above to ensure support for the latest filesystem operation APIs⁸.

Global Installation

OpenSpec is typically installed globally via npm for invocation from any project path:

```
npm install -g @fission-ai/openspec@latest
```

After installation, verify integrity via `openspec --version`⁴.

Interactive Initialization

After entering the project root, run the initialization command to start the configuration wizard:

```
openspec init
```

This process is context-aware. OpenSpec guides users to select the primary AI assistance tool their development team uses. Supported options include but aren’t limited to:

- **Claude Code:** Anthropic’s command-line agent.
- **Cursor:** AI-integrated IDE.
- **GitHub Copilot:** Widely-used code completion tool.

instructions.

Configuration Generation Logic

If users select "Cursor" or "Claude Code," OpenSpec automatically configures tool-specific Slash Commands (like `/openspec;proposal`). For example, for GitHub Copilot CLI, it generates specific prompt templates under `.github/prompts/`; for Windsurf, it generates workflow definitions under `.windsurf/workflows/`⁸.

Non-Interactive Initialization (CI/CD Friendly)

For automation scripts or batch deployments, use parameters to skip the wizard:

```
openspec init --tools claude,cursor
```

This makes OpenSpec easy to integrate into enterprise-grade project scaffolding⁹.

Key Configuration File Details

After initialization, the two most critical static config files in the project are `openspec/project.md` and `openspec/AGENTS.md`.

1. Global Context Anchor: `project.md`

This file isn't just a project introduction—it's the "worldview" through which AI understands the entire engineering effort. Developers should exhaustively describe:

- **Tech Stack:** Specify versions (like TypeScript 5.0, React 18) to prevent AI from using outdated or incompatible syntax.
- **Architecture Patterns:** E.g., "All database access must go through Repository layer; Controller direct queries are strictly forbidden."
- **Code Standards:** Naming conventions, directory structure preferences.
- **Business Domain Knowledge:** Industry-specific terminology and rules.

This preset context greatly reduces AI "hallucinations," making generated code more aligned with project conventions⁶.

2. Agent Behavioral Guidelines: `AGENTS.md`

This file contains a special XML-style markup block `<openspec-instructions>...</openspec-instructions>` maintained by the `openspec update` command. It instructs AI: "When user requests contain words like 'plan,' 'proposal,' or 'spec,' must first read this file"⁶. This mechanism essentially injects a low-level system instruction for AI, forcing it into OpenSpec's working mode.

OpenSpec Workflow Complete Analysis

OpenSpec restructures the development process into a closed-loop engineering cycle. This section details each phase's operational specifics and underlying logic.

Phase One: Proposal & Planning

In traditional development, developers might directly tell AI: "Help me add login functionality with two-factor authentication." In the OpenSpec workflow, this request transforms into a structured proposal process.

Trigger Method

In Cursor or Claude Code, type:

Note: Different tools may adjust the trigger prefix; Kilo Code uses `/openspec-proposal.md`¹⁰.

AI's Reasoning Process:

1. **Context Retrieval:** AI reads `project.md` and related files under `specs/`, analyzing the existing system's authentication module status.
2. **Scaffolding Generation:** AI creates a new directory named `add-user-auth` under `openspec/changes/`.
3. **Document Generation:**
 - **proposal.md:** Articulates change motivation, scope, and expected impact—equivalent to a simplified PRD³.
 - **tasks.md:** This is AI's "action guide" for itself, breaking complex tasks into indivisible atomic steps. E.g., "1.1 Add OTP secret key field to user table," "2.1 Implement TOTP generation algorithm," "3.1 Update frontend login form UI." This breakdown is crucial—it allows AI to "check off" each completed step during subsequent implementation, maintaining progress control⁸.
 - **design.md** (optional): Records technical decisions like which encryption library to choose, database indexing strategies, etc.²

Phase Two: Spec Definition & Deltas

This is the watershed separating OpenSpec from ordinary task management tools. At this stage, AI must explicitly indicate how it intends to modify system behavior. OpenSpec adopts a "delta" mode to describe requirement changes.

Delta Types:

- **ADDED:** Brand new functional requirements.
- **MODIFIED:** Modifications to existing logic. Must include complete revised requirement text.
- **REMOVED:** Deprecated features.

Example: Two-Factor Authentication Spec Delta

AI generates content in `changes/add-user-auth/specs/auth/spec.md` like:

```
## ADDED Requirements

### Requirement: Two-Factor Authentication
Users must provide a second verification factor when logging in.

#### Scenario: OTP Verification Required
- **GIVEN** User has 2FA enabled
- **WHEN** Correct username and password provided
- **THEN** System should return OTP challenge request rather than directly issuing token
```

This structured GIVEN/WHEN/THEN format (similar to Cucumber/Gherkin) makes requirements both clear and testable⁶.

Validation

Before entering coding, developers run the validation command:

```
openspec validate add-user-auth
```

checking before code compilation, intercepting ambiguous requirements .

Phase Three: Implementation & Coding (Apply)

When specs pass both human review and machine validation, actual coding begins.

Execute Command:

```
/openspec:apply add-user-auth
```

AI's Execution Logic:

1. **Read Tasks:** Load tasks.md.
2. **Reference Specs:** AI treats proposal.md and deltas in specs/ as inviolable instructions.
3. **Sequential Progress:** AI modifies source code files sequentially per task list order. It doesn't "wing it" but strictly implements behaviors defined in specs.
4. **State Sync:** After completing each task, AI may update status markers in tasks.md⁷.

Since AI only needs to focus on "how to implement" (How) rather than "what to implement" (What) at this stage, its cognitive load drops dramatically, significantly improving code generation accuracy.

Phase Four: Archive & Merge

After feature development completes and passes tests, the change lifecycle reaches its conclusion.

Execute Command:

```
/openspec:archive add-user-auth --yes
```

System Actions:

1. **Spec Merge:** CLI tool intelligently merges spec deltas from changes/ directory into main specs/ directory. New requirements are appended; modified requirements overwrite old versions.
2. **Cleanup:** The add-user-auth directory is removed or moved to archive.
3. **Single Source Update:** At this point, the specs/ directory accurately reflects the system's latest state including new functionality, ready for the next iteration⁴.

AI Agent Integration & Tool Ecosystem

OpenSpec's power lies in its cross-platform compatibility. It seamlessly integrates with mainstream AI programming tools through specific adaptation layers.

Cursor Integration

Cursor is currently one of OpenSpec's best-supported IDEs.

- **Principle:** OpenSpec leverages Cursor's custom Slash Command functionality.
- **Experience:** Users don't need to leave the editor—directly type `/openspec:proposal` in the Chat window to invoke workflows. OpenSpec-generated files (like proposal.md) open directly in the editor for user review and modification.

summary .

Claude Code Integration

Claude Code provides powerful reasoning capabilities, particularly suitable for handling complex refactoring tasks.

- **Principle:** OpenSpec registers custom commands by configuring Claude Code's `config.toml` or similar mechanisms.
- **Characteristics:** When executing `/openspec:proposal`, Claude Code exhibits extremely strong logical planning capabilities. It can recursively read existing code, with generated `design.md` often containing profound architectural insights¹².
- **Note:** When using Claude Code, sometimes encountering "Plan Mode" restrictions prevents direct file operations (like archiving). In such cases, explicit authorization or manually running CLI commands is needed¹³.

GitHub Copilot CLI Integration

- **Principle:** OpenSpec generates `.prompt.md` files under `.github/prompts/`.
- **Usage:** Copilot CLI automatically recognizes these prompt files, transforming them into slash commands. This makes OpenSpec workflows seamlessly embed into GitHub's native ecosystem, convenient for enterprise user adoption¹⁴.

Windsurf & Kilo Code

These emerging tools typically have automatic workflow discovery capabilities.

- **Configuration:** `openspec init` places definition files in directories like `.kilocode/workflows/`.
- **Automation:** In these tools, `/openspec-apply` can trigger more advanced agent behaviors, even configured to auto-run unit tests after code generation, marking tasks complete only when tests pass⁸.

Advanced Application Patterns & Best Practices

For large teams or legacy projects, OpenSpec offers advanced patterns to handle complex scenarios.

"Retrofitting" Mode

Facing undocumented legacy code, OpenSpec can serve as a reverse engineering tool.

- **Operation:** Create a change proposal named `baseline`.
- **Instruction:** Prompt AI to "read source code in `src/legacy-module` and reverse-generate OpenSpec spec files describing its current behavior."
- **Value:** This not only generates documentation but provides a baseline for refactoring. In subsequent refactoring, you can have AI ensure new implementations still conform to these baseline specs, guaranteeing behavioral consistency¹¹.

Token Efficiency Optimization

In large projects, directly feeding all code to AI instantly exhausts tokens. OpenSpec achieves "load-on-demand" through structured files:

- **Effect:** This streamlined context input enables AI to focus on current slices, significantly reducing token consumption while improving response speed and accuracy².

Team Collaboration & CI/CD Integration

OpenSpec's filesystem architecture naturally supports Git workflows.

- **Code Review:** When submitting PRs, reviewers see not just code (diff) but also proposal.md and specs/ under changes/. This lets reviewers first judge "whether intent is correct" before seeing "whether implementation is correct."
- **Automated Validation:** Add `openspec validate` steps in CI pipelines, mandating all submitted spec files conform to syntax norms, preventing low-quality requirement documents from entering the repo¹².

Troubleshooting & Common Issues Guide

Despite OpenSpec's rigorous workflows, problems can still arise in practice. Here are diagnostic and repair solutions for common failures.

Problem Phenomenon	Possible Cause	Solution
AI ignores specs, writes code directly	Context overload or AGENTS.md not read	1. Run <code>openspec update</code> to refresh instructions. 2. Explicitly require in Prompt: "Read @openspec/AGENTS.md first" ⁶
Command <code>openspec</code> not found	npm global path not in PATH environment variable	Check <code>npm list -g --depth=0</code> , add Node bin directory to Shell config file (like <code>.zshrc</code>) ⁴
Cannot archive (Plan Mode)	AI agent in security-restricted mode, prohibiting file deletion operations	1. Try authorizing in Chat. 2. Downgrade to manually running <code>openspec archive <id></code> in terminal ¹³
Spec validation fails	Spec file missing Scenario or Header format errors	Run <code>openspec validate</code> to view specific errors, ensure each Requirement contains at least one Scenario ¹¹
AI hallucinates/references nonexistent libraries	project.md doesn't clarify tech stack versions	Update project.md, explicitly specify dependency library versions (like "Use React 18, not 16") ⁶

Competitive Analysis: OpenSpec vs. Speckit vs. Kiro

In the AI-assisted development field, OpenSpec isn't the only option. Comparison can more clearly position its applicable scenarios.

Core Scenario	Brownfield development (1→n): Excels at iterating on existing codebases	Greenfield development (0→1): Excels at building new apps from scratch	Fully integrated IDE experience
Change Management Mechanism	Centralized: All files for each change in one independent changes/ folder	Distributed: Updates scattered across multiple spec files	UI-based management, file invisibility high
Lightweight Level	Extremely High: Pure CLI, Markdown-based, no server dependencies	Medium: Generates numerous scaffolding files	Low: Requires specific IDE environment or SaaS account
Cost & Privacy	Free/Open Source: No API keys needed, data localized	Free/Open Source	Some features paid/require cloud sync
Design Philosophy	Structure First	Spec First	Visual SDD

Deep Comparative Analysis

OpenSpec’s biggest advantage is its “change isolation” mechanism. SpecKit tends to directly modify main spec files, easily causing conflicts during multi-person collaboration. OpenSpec’s changes/ directory design makes each feature development like an independent mini-project, merging only at the final moment. This design better aligns with modern software engineering’s Git Flow or Trunk Based Development patterns⁷.

Conclusion & Outlook

OpenSpec represents an important maturation evolution in AI-assisted software development. It profoundly recognizes that while LLMs possess powerful code generation capabilities, they lack inherent mechanisms for maintaining long-term architectural consistency and business logic integrity. By introducing lightweight yet strict file structures and state machine workflows (Proposal -> Apply -> Archive), OpenSpec successfully transforms “vibe coding’s” randomness into repeatable, auditable engineering processes.

For teams dedicated to deeply integrating AI into core development workflows, OpenSpec provides a low-cost, high-return solution. It not only solves token context limits and hallucination problems but, more importantly, re-establishes “specs” core position in development—code is merely implementation details, while specs are the system’s soul. As AI agent capabilities further advance, frameworks like OpenSpec that bridge human intent and machine execution will inevitably become future software industry infrastructure.

References

1. [OpenSpec: NEW Toolkit Ends Vibe Coding! 100x Better Than Vibe Coding \(Full Tutorial\) - YouTube](#)  ²

3. OpenSpec - Lightweight & portable spec driven framework for AI coding assistants! [🔗](#) [🔗](#)² [🔗](#)³ [🔗](#)⁴ [🔗](#)⁵
 4. @fission-ai/openspec - npm [🔗](#) [🔗](#)² [🔗](#)³ [🔗](#)⁴ [🔗](#)⁵ [🔗](#)⁶ [🔗](#)⁷
 5. AGENTS.md [🔗](#)
 6. How to make AI follow your instructions more for free (OpenSpec) [🔗](#) [🔗](#)²
[🔗](#)³ [🔗](#)⁴ [🔗](#)⁵ [🔗](#)⁶ [🔗](#)⁷
 7. From Vibe Coding to Spec-Driven Development: Building Software That Scales
[🔗](#) [🔗](#)² [🔗](#)³ [🔗](#)⁴
 8. Fission-AI/OpenSpec: Spec-driven development (SDD) for AI coding assistants
[🔗](#) [🔗](#)² [🔗](#)³ [🔗](#)⁴
 9. Issue #385 · Fission-AI/OpenSpec - CLI Options [🔗](#)
 10. OpenSpec vs Spec Kit: Choosing the Right AI-Driven Development Workflow
[🔗](#)
 11. openspec-retrofit.md - GitHub Gist [🔗](#) [🔗](#)² [🔗](#)³
 12. A Practical Development Guide Based on OpenSpec + Claude CLI [🔗](#) [🔗](#)²
 13. Mwahahahaha! It lives! · Issue #391 · Fission-AI/OpenSpec [🔗](#) [🔗](#)²
 14. Feature Request: Support custom slash commands from .github/prompts directory [🔗](#)
-

Related Notes

GitHub Spec Kit Deep Dive: AI-Driven Specification Development Methodology

January 8, 2026

An in-depth analysis of GitHub Spec Kit's architecture, workflows, and enterprise applications exploring how Spec-Driven Development solves context loss in AI programming

BMAD-METHOD Guide: Breakthrough Agile AI-Driven Development

January 10, 2026

A comprehensive analysis of BMAD-METHOD's core architecture, deployment, agent roles, and four-phase agile methodology for mastering spec-driven AI development

Notes on 'Deep Work' by Cal Newport

January 10, 2025

Key insights and takeaways from Cal Newport's book on focused work

© 2026. All Rights Reserved.

[Sitemap](#)

[RSS Feed \(English\)](#) [RSS Feed \(中文\)](#) [RSS Feed \(日本語\)](#)