Building Pragmatic AI Agents That Use Tools and APIs

A Unified Guide to DSPy, OpenAI Agents SDK, Claude Agent SDK, Google ADK, and beyond.

(Steve T.)

Building Pragmatic AI Agents That Use Tools and APIs

Building Systems That Use Tools and APIs with DSPy, Pydantic AI, Claude SDK, OpenAI Agents SDK, and Google ADK

Steve T. Team Publications

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

This version was published on 2026-07-03

Leanpub

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 Unified Guide to DSPy, OpenAI Agents SDK, Claude Agent SDK, Google ADK, and Beyond

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

About This Book

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

Copyright and Dedication

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

Chapter 1: Introduction: The Agent Revolution: Why Tools Change Everything

A Tuesday Morning in September 2025

It started with a customer service ticket that nobody wanted to answer.

On an unremarkable Tuesday morning, an airline's support team received a request: "*Please help me book a flight from SFO to JFK on September 1st, my name is Adam.*" The old way would have routed this to a human agent who'd log into the reservation system, search for flights, pick the best one, and confirm the booking: a process taking minutes of human attention and tens of dollars in operational cost.

The new approach used an AI agent armed with six tools: `fetch_flight_info`, `pick_flight`, `book_flight`, `get_user_info`, `cancel_itinerary`, and `file_ticket`. The model didn't hallucinate the answer. It reasoned through the steps (fetch available flights, pick the cheapest option, look up Adam's profile from the user database, book the flight) and returned a confirmation number: `94ldyhsx`. The entire process ran in under five seconds [1,2].

On paper, this is beautiful. In practice, the same agent will stumble when asked to cancel a reservation with an invalid confirmation number, loop endlessly trying to fetch a non-existent itinerary, or silently book the wrong flight because its tool arguments were malformed. The gap between a demo and production is where most AI agent projects die.

This book exists to close that gap.

What "Agentic AI" Actually Means (And What It Doesn't)

The term "agentic AI" has become one of the most overused phrases in technology. Marketing decks show animated diagrams of agents that "think," "plan," and "act", as if the model were a little person inside the computer making conscious decisions. The reality is simpler and more interesting.

An AI agent, at its engineering level, is a system that:

1. Receives a task description from a user
2. Uses an LLM to reason about what information or actions are needed
3. Calls tools to gather that information or take actions in the world
4. Incorporates tool results and repeats until the task is complete

This pattern was formalized in the ReAct (Reason + Act) paper by Yao et al., researchers from Princeton University and Google Research, which showed that interleaving language model reasoning with tool execution produces significantly better results than either approach alone on complex tasks [58]. The paper's insight was deceptively simple: when a model is forced to articulate its reasoning before calling a tool, it makes fewer mistakes. The reasoning step acts as a self-correction mechanism: the model spots its own errors in the thought trace before committing to an action.

But here's what the ReAct paper didn't address, and what every production engineer has learned through painful experience: **the quality of the tools you give the agent matters more than the quality of the model**.

The Single Biggest Factor in Agent Success

Consider two agents. Agent A uses a state-of-the-art model (illustrative example: "GPT-5": a forward-looking projection) with four poorly documented tools that accept generic `dict` arguments. Agent B uses an older model with five precisely described tools that have strict type schemas, informative error messages, and structured outputs. In practice, Agent B will outperform Agent A on the vast majority of real-world tasks.

This comparison illustrates a principle that holds across model generations: tool quality dominates model quality in production settings. Every subsequent framework chapter uses this same pattern: illustrating capabilities with illustrative model names while emphasizing that the architectural patterns apply regardless of which specific model version you deploy.

Why? Because the LLM sees exactly what you define in the tool schema and description before it makes a single tool call. The model uses these descriptions as its primary signal for understanding what each tool does, when to call it, and what to expect. A vague tool description like `"Fetch flight information"` gives the model almost nothing to work with. A rich description like "*Fetch available flights from origin to destination on a specific date. Returns a list of Flight objects with flight_id, date_time, duration, and price. Only call this when the user asks about flight availability for a specific route and date*" gives the model actionable guidance that dramatically reduces hallucinated arguments and missed calls.

This is why the chapters ahead are structured to first establish foundational principles (the anatomy of an agent, tool design), then dive deep into each framework with complete, runnable examples, and finally synthesize cross-cutting patterns for production systems.

Why Tools Change Everything

Before function calling existed, LLMs were generators: they produced text based on patterns in their training data. If you asked an LLM for today's weather in Paris, it would make its best guess based on what it had seen before. It couldn't look up the current temperature because it had no connection to the world beyond its parameters. The model might say "It's probably around 20°C", which sounds confident but is almost certainly wrong.

Function calling, the ability for a model to output structured data that a program can execute, turned LLMs from isolated generators into orchestrators. The model became a reasoning engine; the tools became its hands and eyes in the real world.

This shift has cascading implications for every system you build:

Determinism where there was none. When an agent calls a database query tool, the result is deterministic. A flight search returns exactly what exists. This lets you build verification and validation layers that weren't possible with pure generation. You can write unit tests for your tools independently of the LLM, catch errors before they propagate, and build retry logic that actually works.

Composability. Each tool is a small, testable function. You can unit-test `fetch_flight_info` independently of the LLM. You can swap out weather APIs without changing the agent's reasoning logic. The system decomposes into components you can reason about individually. This is crucial: a welldecomposed tool set means your agent becomes easier to debug, test, and improve over time.

Cost control. A well-designed tool set lets you route simple queries through cheap models (they just need to call a lookup) and reserve expensive reasoning for genuinely complex decisions. This matters enormously at production scale. Consider a customer support agent that handles 10,000 requests per day. If each request involves five tool calls across three model turns, and you're using $5/MTok (input) and $15/MTok (output), you're spending roughly $2,500–$5,000 per day on tokens alone. At that scale, a 30% reduction in tool calls through better tool design translates to hundreds of dollars saved daily.

The Framework Landscape: Five Philosophies

The space has fragmented into distinct philosophies. Each framework represents a different answer to the question: "What's the hardest part of building agents?"

DSPy (from Stanford's NLP group) answers: *writing and maintaining prompts*. DSPy treats agent building as program synthesis: you declare what the system should do with typed signatures, and DSPy optimizes the prompts and few-shot examples automatically. It's a paradigm shift from "write prompts" to "define contracts" [59]. The framework's secret weapon is its optimizer ecosystem: you can define an agent, collect a dataset of input-output pairs, and then run an optimizer to improve the prompts automatically without manually tweaking a single word.

Pydantic AI answers: *type safety and validation*. It brings everything Python developers love about the Pydantic ecosystem (validation, type hints, dependency injection) to agent building. Every tool parameter is validated at runtime by Pydantic schemas. Dependencies flow through a clean injection system. If your team already uses FastAPI and Pydantic, the mental model transfers directly [7]. The framework's underlying `pydantic-graph` finite state machine gives you both flexibility (the model decides within each node) and structural guarantees (the overall graph is fixed).

The Claude Agent SDK answers: *boilerplate*. It takes a minimalist approach: you give Claude a set of tools and permissions, and Claude *executes* them directly: no manual dispatch code needed. It's designed for developers who want Claude to read files, edit code, run commands, and call custom APIs with minimal boilerplate [21]. The SDK runs tools in-process as MCP servers, eliminating subprocess management overhead while maintaining the standard MCP interface.

The OpenAI Agents SDK answers: *orchestration.* It offers three primitives: agents, handoffs, and guardrails. An agent is a model plus instructions and tools. Handoffs let one agent delegate to another: perfect for customer support routing. Guardrails provide input/output validation before and after execution. The framework is lightweight by design, built on the principle that you should be able to build a multi-agent system in a single file [25].

Google's Agent Development Kit (ADK) answers: *deterministic control.* It takes an enterprise-first approach with template workflows: SequentialAgent chains agents step-by-step, ParallelAgent fans out independent tasks concurrently, and LoopAgent implements iterative refinement loops. Starting with version 2.0, ADK added graph-based workflows that let you compose both AI-powered agents and deterministic execution nodes into a flexible execution graph [37]. It's designed for multi-agent systems that need predictable execution paths and enterprise deployment (Cloud Run, GKE, Agent Engine).

Each framework has strengths and trade-offs. None is universally superior. The right choice depends on your team's expertise, your deployment environment, and the complexity of the tasks your agents need to handle.

What This Book Will Do

This book will walk you through each framework with complete, runnable code examples. You'll build:

- A customer service agent in DSPy that books flights using a ReAct loop with escalation paths
- A weather agent in Pydantic AI with dependency injection, structured outputs, and usage limits
- A bug-finding agent in the Claude Agent SDK that reads, analyzes, and fixes real source code
- A multi-agent support system in the OpenAI Agents SDK with handoffs, guardrails, and sandbox agents
- An iterative document improvement pipeline in Google ADK using sequential and loop workflows

Beyond the code, you'll learn cross-cutting patterns: how to design tools that agents actually use well, how to handle errors gracefully, how to test agent behavior systematically, how to optimize costs in production, and how to choose the right framework for your specific context.

The Model Context Protocol (MCP) (an open standard for connecting AI agents to external systems with over 10,000 public servers and 97 million monthly SDK downloads as of early 2026 [48]) will appear throughout as the common layer. Understanding MCP isn't required to use any of these frameworks, but it's becoming essential infrastructure. All five frameworks we cover support MCP either natively or through adapters.

What This Book Won't Do

This book does not cover:

- Fine-tuning LLMs from scratch (that's a different discipline, though we touch on DSPy's fine-tuning capabilities)
- Building RAG systems from scratch (though we discuss tool-based retrieval as an alternative)
- Multi-agent communication protocols beyond what each framework provides natively (we cover A2A briefly in the ADK chapter and the future chapter)
- Hardware or infrastructure for hosting agents at scale (deployment patterns are covered, but not Kubernetes configuration)

If you're looking for those topics, they deserve their own books. This one stays focused on the core engineering challenge: making AI agents that use tools and APIs reliably and pragmatically.

How to Read This Book

This book is designed for two reading styles:

Linear reading (recommended for first-time readers): Work through the chapters in order. The early chapters establish foundational concepts (the agent loop, tool design) that recur throughout the framework-specific chapters. Reading linearly ensures you understand these concepts before diving into framework-specific details.

Reference reading (recommended for experienced engineers): Jump to the framework chapter relevant to your current project. Each framework chapter is self-contained enough to be read independently, with references back to earlier chapters for foundational concepts. The cross-framework patterns chapter (Chapter 9) is useful at any point for comparing approaches.

Each framework chapter follows a consistent structure:

1. **Philosophy**: What problem does this framework solve?
2. **Complete working example**: A full code walkthrough
3. **Deep dives**: Feature-by-feature exploration with practical examples
4. **War stories**: Composite scenarios for pedagogical illustration (see disclaimer above)
5. **When to use / when not to use**: Honest assessment of trade-offs

The cross-framework patterns chapter (Chapter 9) synthesizes what we've learned into actionable principles that apply regardless of which framework you choose. The production chapter (Chapter 10) covers testing, cost management, deployment, and monitoring: the stuff that separates a demo from a system your customers actually depend on.

The code examples are drawn from official documentation and tutorials [1,7,21,25,34], verified against the latest API surfaces as of mid-2026. If you want to run the examples yourself, all frameworks are open-source and available on GitHub.

Chapter 2: The Anatomy of an AI Agent

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

The Agent Loop: Observe, Reason, Act

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

A Concrete Example: The Airline Agent

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

Function Calling: The Foundational Primitive

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

The Spectrum of Autonomy

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

Deterministic vs. Probabilistic Control Flow

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

Context Window Management: The Hidden Bottleneck

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

A Composite Scenario: The Infinite Loop

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

The MCP Layer: The Universal Language

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

MCP Security: The Hidden Risk Surface

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

Chapter Summary

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

Chapter 3: Designing Tools That Agents Can Use Well

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

The Anatomy of a Good Tool Definition

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

1. Name: Be Specific, Not Clever

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

2. Description: The Most Important Field

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

3. Input Schema: Strict Types, Clear Constraints

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

Error Handling: Making Failures Informative

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

Structured Outputs: When Tools Return Data, Not Just Text

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

The Context Window Tax of Tool Descriptions

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

Composite Scenario: The 50-Tool Agent (Context Window Overload Pattern)

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

Tool Design Case Studies

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

Case Study 1: Travel Booking Tools

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

Case Study 2: Code Analysis Tools

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

Anti-Patterns to Avoid

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

Chapter Summary

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

Chapter 4: DSPy: Programming LLM Pipelines, Not Prompts

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

The DSPy Philosophy: Signatures Over Prompts

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

Building the Airline Agent: A Complete Walkthrough

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

The ReAct Loop Internals: What DSPy Does Under the Hood

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

Trajectory Inspection: Debugging the Agent's Thinking

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

Escalation: The File-Ticket Pattern

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

Optimization: DSPy's Secret Weapon

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

Composite Scenario: Optimizing a Failing RAG Pipeline (DSPy Optimization Pattern)

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

DSPy's Design Philosophy: Why Signatures Over Prompts?

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

Why DSPy's Optimizers Work: The Mechanics of MIPROv2

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

DSPy vs. Traditional Prompt Engineering: A Side-by-Side Comparison

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

Debugging DSPy Agents: A Walkthrough

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

DSPy in Production: Real-World Deployments

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

DSPy's Limitations: When to Avoid It

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

DSPy vs. Pydantic AI: A Deeper Comparison

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

When to Use DSPy (and When Not To)

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

DSPy and MCP

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

Inference-Time vs. Optimization-Time Trade-offs

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

Chapter Summary

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

Chapter 5: Pydantic AI: Type-Safe Agents the Python Way

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

Core Concepts: Agent as a Typed Container

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

Five Execution Pathways

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

Tools: `@agent.tool` vs `@agent.tool_plain`

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

Dependency Injection via `RunContext`

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

Structured Outputs with Pydantic Models

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

Usage Limits and Cost Control

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

Model Settings and Configuration

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

Concurrency Limiting and Backpressure

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

Streaming Modes: Three Levels of Visibility

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

Self-Correction and Retry Budgets

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

Observability: Logfire and OpenTelemetry

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

Declarative Configuration: YAML Agent Specs

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

Runs vs. Conversations: Message History

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

Durable Execution: Surviving API Failures

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

A Composite Scenario: Migrating from Flask to Pydantic AI

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

Pydantic AI vs. OpenAI Agents SDK: A Deeper Comparison

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

Pydantic AI vs. OpenAI Agents SDK: Guardrails Compared

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

Pydantic AI's Durable Execution: Production-Grade Reliability

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

When to Use Pydantic AI (and When Not To)

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

Chapter Summary

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

Chapter 6: Claude Agent SDK: In-Process Tools and Built-in Execution

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

The `query()` Function: Your Entry Point

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

Built-in Tools: The Complete Toolset

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

Permission Modes: Controlling Autonomy

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

Custom Tools: The `@tool` Decorator and In-Process MCP Servers

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

Error Handling: `isError` vs Exceptions

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

Tool Annotations: Behavioral Metadata

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

Hooks: Intercepting Agent Behavior at Key Points

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

Subagents: Spawning Specialized Agents from Within a Run

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

Sessions and Multi-Turn Conversations

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

MCP Server Integration

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

Third-Party Provider Support

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

SDK vs. Claude Code CLI vs. Managed Agents

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

Claude Agent SDK's In-Process MCP Servers: Why It Matters

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

Claude Agent SDK's Permission Modes: A Deep Dive

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

Claude Agent SDK's Hooks: Deterministic Control Over Probabilistic Behavior

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

When to Use the Claude Agent SDK (and When Not To)

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

Chapter Summary

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

Chapter 7: OpenAI Agents SDK: Lightweight Orchestration with Handoffs

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

The Three Primitives: Agents, Handoffs, Guardrails

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

Function Tools with Automatic Schema Generation

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

Constraining Arguments with Pydantic Field

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

Tool Timeouts and Error Handling

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

Tool Context and Dependencies

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

Hosted Tools: OpenAI's Built-in Capabilities

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

Tool Search for Deferred Loading

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

Agents as Tools: A Second Multi-Agent Pattern

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

Handoffs: The Primary Multi-Agent Pattern

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

Guardrails: Input, Output, and Tool Validation

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

Sandbox Agents: Isolated Execution

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

Tracing and Observability

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

Sessions: Persistent Memory

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

MCP Server Integration

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

Multi-Provider Support

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

OpenAI Agents SDK's Handoff Mechanism: How It Actually Works

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

OpenAI Agents SDK vs. Pydantic AI: Runtime Overhead Comparison

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

OpenAI Agents SDK's Sandbox Agents: Security Through Isolation

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

When to Use the OpenAI Agents SDK (and When Not To)

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

Chapter Summary

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

Chapter 8: Google ADK: Graph-Based Workflows for Enterprise Scale

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

Template Workflow Agents: The Foundation

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

State Management: Shared State Namespace

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

Custom Tools: FunctionTool and Beyond

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

ADK 2.0: Graph-Based Workflows (Workflow Runtime)

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

Dynamic Workflows: Code-Based Logic

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

Collaborative Workflows: Coordinator Agents and Subagents

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

Skills: Loading Domain Expertise on Demand

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

Evaluation and Testing

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

A2A Protocol: Cross-Framework Communication

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

Deployment Options

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

Multi-Language Support

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

ADK for Android

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

Google ADK's Workflow Runtime vs. Template Agents: When to Use Which

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

ADK's Collaborative Workflows: Coordinator Agents and Subagents

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

ADK's Skills Toolset: On-Demand Domain Expertise

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

Google ADK's Multi-Language Support: Why It Matters

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

When to Use Google ADK (and When Not To)

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

Chapter Summary

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

Chapter 9: Cross-Framework Patterns: What Works Everywhere

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

Pattern 1: Tool Design Principles That Are Framework-Agnostic

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

Pattern 2: Error Handling Strategies: Three Levels

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

Pattern 3: Observability and Tracing

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

Pattern 4: Memory and Context Management

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

Pattern 5: Security Considerations

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

Pattern 6: Performance Optimization

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

Pattern 7: The Human-in-the-Loop Spectrum

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

Cross-Framework Migration: What Changes When You Switch

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

Cost Comparison: Token Costs Per Framework/Model

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

Community Health and Ecosystem Maturity

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

Chapter Summary

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

Chapter 10: Productionizing Agent Systems

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

Testing Strategies: A Multi-Layer Approach

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

Cost Management: Seven Strategies

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

Deployment Patterns

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

Monitoring and Observability

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

A/B Testing Agent Configurations

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

Scaling Considerations

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

Composite Scenario: From Prototype to 10k Daily Requests (Production Scaling Pattern)

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

Security Hardening and Compliance

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

CI/CD Integration for Agent Systems

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

Cost/Latency Benchmarking Tables

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

State Management Deep Dive

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

Chapter Summary

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

Chapter 11: The Future of Agent Tool Use

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

MCP Ecosystem Growth: What's Certain

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

Multi-Agent Collaboration Standards: What's Emerging

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

Fine-Tuning vs. Prompting for Tool Use: The Convergence

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

Hardware Acceleration and Low-Latency Agents

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

Speculation: What's Next?

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

The Economic Argument: When Do Agents Pay for Themselves?

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

What's Genuinely Uncertain

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

Chapter Summary

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

Conclusion: Choosing Your Path

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

The Decision Matrix: A Refined Guide

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

The Pragmatic Approach: Seven Principles

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

A Decision Framework: When to Use What

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

The Economics of Agents

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

The Framework Landscape: A Snapshot

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

The Big Picture: Agents as Software Engineering

Glossary of Key Terms

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

Index

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

References

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

Framework Documentation

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

MCP Ecosystem

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

MCP Security

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

Academic Papers

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

Industry Analysis and Additional Resources

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