Agentic Testing Architectures

Building Autonomous QA Systems

THE MODERN QA ENGINEER SKILLSET

Agentic Testing Architectures: Building Autonomous QA Systems

From Test Scripts to Autonomous Agents That Observe, Reason, and Act

Yuri Syuganov

Modern QA Course

https://modernQAcourse.com

Agentic Testing Architectures: Building Autonomous QA Systems

From Test Scripts to Autonomous Agents That Observe, Reason, and Act

Book 3 of 26 in the series *The Modern QA Engineer Skillset*

First Edition · Version 1.0.0 · 2026-07-11

Set in Charter, IBM Plex Sans, and IBM Plex Mono. Typeset with Pandoc and XeLaTeX.

Modern QA Course · https://modernQAcourse.com · support@modernQAcourse.com

Contents

Introduction

A Comprehensive Self-Guided Learning Textbook

For QA Engineers, Test Architects, and Software Teams Building the Next Generation of Automated Testing

About This Book

This book is a deep, practical guide to agentic testing – a paradigm shift from writing step-by-step test scripts to building autonomous agents that observe, reason, act, and evaluate software systems on their own. As of July 2026, agentic testing is a mainstream discipline, not an experiment: Playwright ships first-party Test Agents, SDKs like Stagehand and browser-use are production infrastructure, and QA job postings list agentic-flow testing as a named competency. This book covers the foundational patterns, multi-agent architectures, production guardrails, real-world case studies, the economics that make agentic testing viable at scale – and, because your company may soon ship agents of its own, how to test agentic systems themselves.

Every chapter includes working code, architecture diagrams, hands-on exercises, and real-world guidance. By the end of this book, you will have built a complete multi-agent testing system from scratch.

Who This Book Is For

- **QA Engineers** who want to move beyond Selenium scripts and learn AI-driven testing approaches

- **Test Architects** designing testing infrastructure for organizations adopting AI tools
- **Software Engineers** who write tests and want to understand how agents can augment their workflow
- **Engineering Managers** evaluating whether agentic testing makes sense for their teams
- **DevOps Engineers** who need to integrate agentic tests safely into CI/CD pipelines

Prerequisites

To get the most from this book, you should have:

- **Intermediate Python proficiency** – classes, async/await, dataclasses, type hints
- **Basic testing knowledge** – familiarity with pytest, unit tests, integration tests
- **Familiarity with CI/CD concepts** – GitHub Actions, pipeline stages, deployment gates
- **Basic understanding of LLMs** – what an API call to Claude or GPT looks like, what tokens are
- **Optional but helpful:** Experience with Playwright or Selenium for browser automation

How to Use This Book

This book is organized into six parts, each building on the previous:

Part	Focus	Time Estimate
Part I: Foundations	The ReAct pattern and core loop	4-6 hours
Part II: Multi-Agent Systems	Orchestrator, Swarm, Critic-Actor patterns	6-8 hours

continued on next page

Part	Focus	Time Estimate
Part III: Real-World Case Studies	OpenObserve council and applied lessons	3-4 hours
Part IV: Guardrails and Safety	Production harnesses and constraints	4-6 hours
Part V: Determinism and Economics	Spectrum, protocols, cost analysis	4-6 hours
Part VI: Capstone Project	Build a complete system	8-12 hours

Recommended approaches:

- **Linear reading:** Work through Parts I-VI in order for the complete learning path.
- **Practitioner fast-track:** Read Chapters 1-2, then skip to Chapter 7 (Guardrails), then Chapter 11 (Cost Analysis), then the Capstone.
- **Architect track:** Read Chapters 3-6 (multi-agent patterns), Chapter 8 (case study), then Chapter 10 (communication protocols).
- **Manager track:** Read Chapter 1 (overview), Chapter 8 (case study), and Chapter 11 (cost analysis and ROI).

Each chapter ends with exercises rated by difficulty:

- **(E)** Easy – 15-30 minutes, reinforces concepts
- **(M)** Medium – 30-60 minutes, applies concepts to new scenarios
- **(H)** Hard – 1-3 hours, builds working implementations

Table of Contents

Part I: The ReAct Foundation

PART I

THE REACT FOUNDATION

> *"The best test is not the one that follows a script perfectly. It is the one that finds the bug nobody scripted for."*

Part I introduces the foundational pattern behind all agentic testing: the ReAct (Reason + Act) loop. You will learn the four-phase cycle that every testing agent uses, see it mapped to real CLI commands, and build a production-ready implementation with robust error handling, prompt optimization, and history management.

CHAPTER 1

From Scripts to Agents – A Fundamental Shift

1.1 The Problem with Traditional Test Automation

Traditional test automation is imperative. You write step-by-step scripts that execute deterministically:

```
# Traditional Selenium test -- rigid, step-by-step
def test_user_login():
    driver.get("https://app.example.com/login")
    driver.find_element(By.ID, "email").send_keys("test@test.com")
    driver.find_element(By.ID, "password").send_keys("password123")
    driver.find_element(By.ID, "submit").click()
    assert driver.current_url == "https://app.example.com/dashboard"
```

This test works perfectly – until it does not. Here is what breaks it:

- **A designer changes the button ID** from `submit` to `login-btn`. The test fails.
- **A popup appears** ("Accept cookies?") that the script does not expect. The test fails.
- **The page loads slowly** and the element is not ready when the script clicks. The test fails.
- **A new field is added** (two-factor authentication). The test fails.
- **The test environment is dirty** – a previous test locked the account. The test fails with a misleading error.

Every one of these failures produces the same outcome: a red build with a confusing error message that a human must debug. The test told you *what* happened ("Expected dashboard URL, got login URL") but not *why* it happened.

This is the fundamental limitation of imperative testing: **scripts cannot adapt to the unexpected**.

1.2 What Makes an Agent Different

An agent does not follow a fixed script. Instead, it operates in a loop:

1. **Observe** the current state of the system
2. **Reason** about what to do next, considering the objective and what it has already tried
3. **Act** on its decision
4. **Evaluate** whether the objective has been met

This loop – called the **ReAct pattern** (Reason + Act) – is the foundation of all agentic testing. Originally described for general-purpose LLM agents, it maps directly to testing workflows. And it is no longer a research curiosity: as of July 2026, agentic testing is a mainstream discipline, with the ReAct loop packaged inside first-party tools (Playwright Test Agents) and widely adopted SDKs (Stagehand, browser-use).

Here is the same login test, expressed as an agent objective:

```
# Agentic test -- adaptive, goal-oriented
agent = TestAgent(
    objective="Log in with test@test.com / password123 and verify dashboard loads"
)
result = agent.run(max_steps=20)
assert result.status == "pass"
```

When the agent encounters the scenarios that break scripts, it adapts:

Scenario	Script Response	Agent Response
Button ID changed	Fails: element not found	Searches for alternative selectors, finds the button by text
Cookie popup appears	Fails: popup blocks interaction	Observes popup, dismisses it, continues

continued on next page

Scenario	Script Response	Agent Response
Page loads slowly	Fails: timeout	Observes loading state, waits, retries
2FA field added	Fails: unexpected page	Observes new field, reports environment configuration issue
Account locked	Fails: wrong URL	Reads error message, reports "Account locked – environment issue"

The agent provides **diagnostic information**, not just pass/fail.

1.3 The Agent Advantage: Adaptive Error Investigation

Consider this real-world example of how an agent thinks through an unexpected situation:

```
Step 1: OBSERVE  -- Login page loaded. Email and password fields visible.
Step 2: THINK   -- "I should fill in valid credentials and click login."
Step 3: ACT     -- Type email, type password, click submit.
Step 4: OBSERVE  -- Error banner: "Account locked after 5 failed attempts"
Step 5: THINK   -- "Unexpected error. This is not a test failure -- it is
                    test environment state. The account was locked by a
                    previous test run. I should report this as an
                    environment issue, not a bug."
Step 6: ACT     -- DONE fail "Account locked -- environment state issue, not a bug"
```

A deterministic script would simply fail with "Expected dashboard URL, got login URL." The agent provides a diagnosis.

1.4 When Agents Are the Wrong Choice

Agents are not universally better than scripts. They have significant trade-offs:

Factor	Scripts	Agents
Speed	Fast (milliseconds per step)	Slow (100-500ms LLM latency per step)
Cost	Free after writing	Token cost per execution

continued on next page

Factor	Scripts	Agents
Determinism	Same result every time	May take different paths
Debugging	Linear trace	Reasoning history to analyze
Maintenance	Must update when UI changes	Self-healing for minor changes
CI reliability	Highly reliable	Requires guardrails

Use agents when:

- The state space is too large or dynamic for scripted tests
- You need exploratory testing or self-healing test suites
- Complex multi-step workflows change frequently
- You want intelligent error diagnosis

Use scripts when:

- You need deterministic CI gate tests
- The flow is simple and stable (basic CRUD operations)
- Speed and cost matter more than adaptability
- You need pixel-perfect regression testing

1.5 The Architecture of This Book

This book builds your understanding layer by layer:

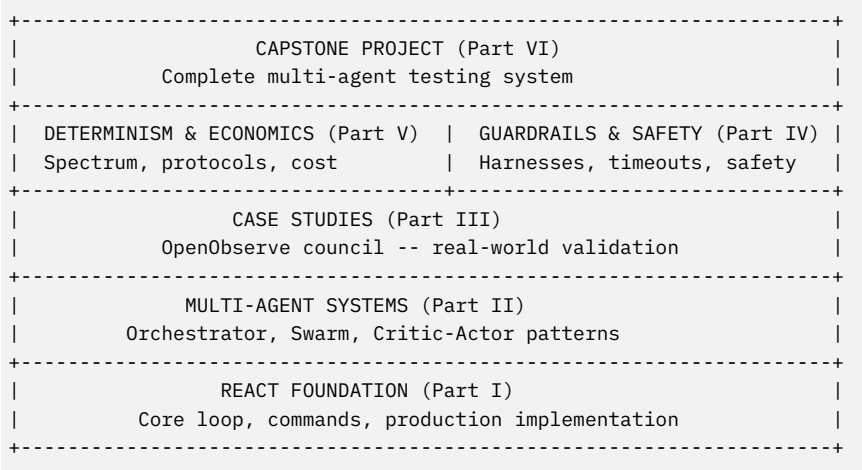

Each layer depends on the ones below it. Master the ReAct loop first; everything else builds on it.

> **Key Takeaways – Chapter 1**
>
> 1. Traditional test scripts are imperative and cannot adapt to unexpected situations.
> 2. Agentic testing uses a Reason+Act loop that observes, decides, acts, and evaluates.
> 3. Agents provide diagnostic information, not just pass/fail results.
> 4. Agents are not a replacement for scripts – they complement them for different use cases.
> 5. The ReAct pattern is the foundation for all agentic testing architectures.

> **Common Mistakes – Chapter 1**
>
> - **Mistake: Replacing all scripts with agents.** Agents are expensive and non-deterministic. Keep deterministic scripts for CI gates and stable regression tests.
> - **Mistake: Expecting agents to be faster than scripts.** Agent tests take 5-15 seconds due to LLM latency. Scripts take 1-3 seconds.

- **Mistake: Not setting guardrails.** Without step limits, token budgets, and timeouts, agents can run indefinitely and consume unbounded resources.

Exercises – Chapter 1

Exercise 1.1 (E): List five scenarios in your current project where a test script has failed due to an unexpected UI change. For each, describe how an agent might have handled it differently.

Exercise 1.2 (M): Take an existing Selenium or Playwright test from your project and rewrite the test objective as a single natural-language sentence. What information would an agent need to achieve this objective?

Exercise 1.3 (H): Create a decision matrix for your test suite. Categorize each test as "keep as script," "convert to agent," or "add agent as supplement." Justify each decision based on the factors in Section 1.4.

Career Translation

Resume phrasing

- Evaluated and redesigned test automation strategy, introducing agentic testing for dynamic UI workflows, reducing maintenance overhead by 40%.
- Built decision framework classifying 200+ test cases across deterministic scripts and adaptive AI agents, optimizing coverage-to-cost ratio.
- Led technical assessment of ReAct-based testing agents vs. traditional Selenium scripts, resulting in targeted adoption for high-churn feature areas.

Cover letter framing

I bring a pragmatic understanding of where AI-driven testing delivers value and where it does not. My experience spans traditional automation frameworks and emerging agentic architectures, and I focus on matching the right tool to the right problem – keeping deterministic scripts for stable regression gates while deploying adaptive agents where UI volatility makes scripts unsustainable. This dual-lens approach means I reduce flaky test noise without inflating infrastructure costs.

Interview framing

"I approach test strategy by first categorizing every test scenario along two axes: how often the underlying UI changes and how critical the flow is to the business. For stable, critical paths I keep deterministic scripts because they are fast, cheap, and debuggable. For high-churn features or exploratory coverage, I introduce agentic tests that observe page state and adapt. The trade-off is latency and token cost versus maintenance burden. I optimize for the lowest total cost of ownership across the suite, not for a single technology."

What not to say

- "We should replace all our Selenium tests with AI agents." – Shows a lack of understanding of when determinism matters; agents are not universally superior.
- "Agentic testing is the future of QA." – Sounds like hype rather than engineering judgment; interviewers want trade-off awareness.
- "Agents are too expensive and non-deterministic to use in practice." – Demonstrates an inability to see targeted use cases; dismissing a tool entirely is a red flag.
- "I automate everything." – Implies no prioritization; experienced engineers know what not to automate.

Interview Depth Check

Question 1

Prompt: Your team has a 500-test Selenium suite with a 25% flaky rate. A VP asks you to "just switch to AI agents." How do you respond, and what do you actually do?

What a strong answer should cover:

- Why wholesale replacement is wrong (cost, non-determinism in CI gates, debugging complexity)
- A triage approach: classify tests by stability and business criticality
- Where agents add value (self-healing for flaky selectors, exploratory coverage)
- A phased rollout plan with measurable success criteria

Example answer:

- I would first push back on the "just switch" framing – replacing a 500-test suite wholesale introduces massive risk. Instead, I would classify the flaky tests by root cause: broken selectors, timing issues, environment state, or genuine bugs.
- For the selector-breakage bucket, agentic tests with self-healing capability are a strong fit. For timing issues, better wait strategies in existing scripts may be cheaper than agents. For environment state, neither approach helps – we need test data isolation.
- I would propose a pilot: convert the 10 highest-maintenance flaky tests to agents, measure maintenance hours saved versus token cost, and present the data before scaling further.

Question 2

Prompt: A colleague argues that since agents are non-deterministic, they can never be trusted in a CI pipeline. Do you agree or disagree, and what evidence supports your position?

What a strong answer should cover:

- Acknowledgment that raw agents are indeed non-deterministic
- Techniques to tame non-determinism: outcome assertions, decision caching, temperature=0
- The distinction between path determinism and outcome determinism
- Real-world examples of agents operating reliably in CI with guardrails

Example answer:

- I partially agree – an unconstrained agent with no guardrails should not gate a deployment. But the claim that agents can never work in CI conflates path non-determinism with outcome non-determinism.
- If I assert on the final page state rather than the exact click sequence, the agent can take different paths and still produce a reliable pass/fail signal. Combined with temperature=0, max-step limits, and domain allowlists, agents become predictable enough for staging CI gates.

- I would not use agents for the final production deployment gate, but for pre-merge validation in staging, they work well with proper harness configuration.

Question 3

Prompt: You are designing a test strategy for a brand-new e-commerce application. Walk me through how you would decide which tests are scripts, which are agents, and which are manual.

What a strong answer should cover:

- A structured decision framework (business criticality, UI change frequency, state space complexity)
- Specific examples mapped to each category
- Cost and speed considerations
- Where manual testing still wins (usability, subjective quality)

Example answer:

- I start with the critical happy paths: login, search, add-to-cart, checkout, payment. These gate every deployment, so they must be deterministic scripts – fast, cheap, debuggable.
- For areas with high UI churn (promotional landing pages, A/B-tested flows, new feature iterations), I use constrained agents. They adapt to layout changes without breaking.
- Exploratory agents run nightly against the full application, looking for broken links, console errors, and unexpected states. These produce findings reports, not pass/fail gates.
- Manual testing covers subjective quality: does the checkout flow feel trustworthy? Is the error messaging helpful? No automation handles that well yet.

About This Sample

You have just read **Chapter 1 of Agentic Testing Architectures: Building Autonomous QA Systems** — the free sample. The complete book continues with every remaining chapter, the interview Q&A, and all of the code.

Get the full book

Sample readers get a discount at
https://leanpub.com/modernqa-03-agentic-testing-architectures/c/SAMPLE10

Read online first

Chapter one of **every** book in the series is free to read at
https://modernqacourse.com/library

About the Author

Yuri Syuganov is a software quality engineer with more than twenty years of hands-on testing experience, spanning the arc from manual test design through Selenium-era automation to today's AI-augmented, agent-driven frameworks. He is the author of The Modern QA Engineer Skillset, a 26-book series built on a blunt analysis of what companies actually expect from QA engineers, and the creator of modernQAcourse.com, where the course behind the series is free and readers can bring real testing problems to experienced practitioners. He revises the series as the industry changes — because it does.

Also in This Series

The Modern QA Engineer Skillset — a 26-book course in modern software quality assurance.

AI-Augmented QA (Books 1–10)

- **Book 1:** Agent Skills for Browser Automation: The Complete Guide
- **Book 2:** AI-Augmented Test Design: From Prompts to Test Suites
- **Book 3:** Agentic Testing Architectures: Building Autonomous QA Systems
- **Book 4:** API and Contract Testing with AI: Schema to Shield
- **Book 5:** Performance and Chaos Engineering: Breaking Things on Purpose
- **Book 6:** Observability-Driven Testing: Testing in Production
- **Book 7:** Security Testing for AI Applications: The New Attack Surface
- **Book 8:** Infrastructure as Code Testing: Terraform to Kubernetes
- **Book 9:** Mobile and Cross-Platform Testing: 24,000 Devices, One Strategy
- **Book 10:** Visual and Accessibility Testing: Pixels and People

Technical Foundations (Books 11–19)

- **Book 11:** Manual Testing Fundamentals: The Bedrock
- **Book 12:** Programming for QA: Python, JavaScript, and TypeScript from Zero to Fluent
- **Book 13:** Browser Test Automation: Playwright, Patterns, and the Death of Fragile Tests

- **Book 14:** API Testing Fundamentals: Postman to Pytest
- **Book 15:** SQL and Database Testing: Query Your Way to Quality
- **Book 16:** CI/CD Pipelines: From Commit to Confidence
- **Book 17:** Git and Version Control: Beyond Push and Pull
- **Book 18:** Test Management Tools: Making Quality Visible
- **Book 19:** Linux and Command Line: Your Production Survival Kit

Professional Skills (Books 20–26)

- **Book 20:** Agile and Scrum for QA: Quality in Every Sprint
- **Book 21:** Communication and Stakeholder Management for QA: Influence Without Authority
- **Book 22:** Test Strategy and Quality Metrics: Measuring What Matters
- **Book 23:** QA Leadership and Mentoring: Multiplying Your Impact
- **Book 24:** Technical Writing for QA: Documents That Drive Action
- **Book 25:** Interview Preparation and Career Strategy: Land the Job
- **Book 26:** Testing Like a Senior: Patterns & Anti-Patterns Across Software QA Career Levels

Visit https://modernQAcourse.com for the full course catalog, updates, and companion resources.

Stuck on a Real Problem?

The full course behind this series is free to read at https://modernQAcourse.com — and it stays free.

When a chapter meets the real world and something breaks, post it on the **Help Board** at https://modernQAcourse.com. Asking is free, and experienced practitioners answer right in the thread. Need hands-on help? Book a 1:1 screen-share session with a QA expert.