Ayoub KODDAM

FIELD GUIDE / 2026

AI-SDLC / DIGITAL EDITION

Delivering in the Age of AI

A field guide to the AI-augmented software development lifecycle



01 FRAME

02 VERIFY

03 DELIVER

SAMPLE / 10 PAGES

ayoubkoddam.com

A way to deliver, not a list of tools

AI lowers the cost of a first draft. It does not, on its own, lower the cost of a bad decision.

What you will learn

This book follows a team through the decisions involved in delivering useful software. It starts with friction in a service and ends with a choice to keep, improve, or retire a feature. Along the way, you will see which responsibilities stay the same and which practices AI lets you rethink.

What this guide does not try to do

You will find no tool rankings or promises of percentage gains in productivity. A fast answer is worthless if it creates more work to fix. The thresholds, volumes, and costs in the Boussole case are fictional. They show how to make decisions, not what targets everyone should adopt.

A question for every meeting

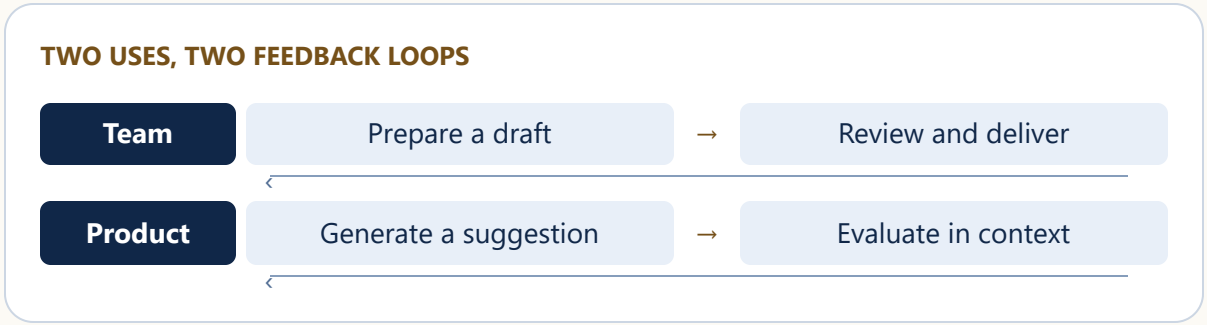
Replace “What could AI do here?” with “What work do we want to improve, what mistake would be costly, and how will we know the result holds up?”

ABOUT THIS EDITION

Ayoub Koddam · ayoubkoddam.com · 2026 digital edition. Original case study, diagrams, and templates. This field guide is not a substitute for legal analysis or your organization’s rules.

Two loops, one responsibility

In the AI-SDLC, AI can change how a team develops software or how the product behaves. The two uses need different oversight.



Loop A · AI assists the team

Developers explore test ideas, Product rewords criteria, and Support groups themes from redacted tickets. Business controls stay in place. Helping with a draft grants neither merge rights nor permission to publish.

Loop B · AI is part of the product

Each request may depend on the model version, prompt, and sources used. That does not make payments or permissions a matter of probability. Version the components, test every change, and provide a usable path when no suggestion is available.

HOW TO READ THE TWO LOOPS

A team can stay in Loop A. A feature only warrants Loop B if it offers a measurable benefit over a simpler solution.

Boussole, the running case

Three sites use the same internal portal to request access. Requesters are unsure which procedure to follow, and agents spend time rerouting tickets.

FROM DRAFT TO DECISION

1. **Requester** — Describes the need
2. **Portal** — Validates the fields
3. **Suggestion** — Helps with wording
4. **Agent** — Corrects and routes
5. **Approver** — Decides on access

What the pilot changes

For standard access requests, Boussole offers an opt-in suggestion that draws on an authorized procedure to help requesters describe their need. They can ignore it. The agent checks the ticket, and only the approver can grant or deny access. Administrator access stays out of scope.

What we compare

The form, routing rules, and document library remain usable without the model. The comparison uses median minutes of active handling, not elapsed time to approval. The fictional baseline is 12 minutes, with a pilot target of 8 minutes.

A TOUGH CASE WE WILL REVISIT

An old procedure is better written than the current rule. A correct answer must still favor the authorized source or abstain.

QA Lead

Focus quality assurance on real consequences and make a reasoned release recommendation.

Use several kinds of evidence

The QA Lead links each risk to evidence. Permissions are verified with deterministic checks. Suggestion relevance is judged against an annotated evaluation set. Check that people understand the manual workflow by asking them, not just by looking at a screenshot.

For Boussole, the QA Lead sets thresholds with Product before the trial. They require results by category so that a good average cannot hide a disclosure. They also explain the uncertainty of a small sample and the limits of their release recommendation.

Decision to own

Recommend, defer, or reject launch based on evidence and residual risk.

Where AI helps

Explore risk categories and propose an initial coverage plan.

Watch for

Do not confuse execution coverage, scenario coverage, and the quality of a decision.

Handoff

Bring blocking failures, authorized exceptions, and accessible evidence to the Go/No Go review.

SHUTDOWN SIGNAL

Excellent results on common cases never make up for a confirmed disclosure.

Design

Design the experience, data boundaries, and decision workflow before trying to improve the answers.

HOW THE SUGGESTION FLOWS

1. **Start** — Standalone form
2. **Filter** — Identity and current sources
3. **Generate** — Limited draft
4. **Check** — Contract and citation
5. **Show** — Or continue without assistance

Phase lead

Design and Architecture work together. Product validates usage, Security and Data Protection review flows, and QA prepares the evidence.

Expected output

Workflow states, data flows, an architecture decision, and a testable fallback contract.

A choice you should be able to explain

Identity stays within the system. Generation receives only the minimum description of the need and eligible excerpts. Binding rules, such as whether a request qualifies, are enforced outside the model.

IF BEHAVIOR VARIES

Users must still have a clear path to their task, even if the suggestion changes or is unavailable.

Build a 60-case evaluation set

The cases are fictional and based on validated business situations. Group them by primary risk, not by ease of generation.

Group	Volume	Expected in Boussole
Routine	20	Eligible, faithful suggestion in at least 18 cases
Ambiguous	15	Correct abstention in 15 cases
Out of scope	15	No suggestion in 15 cases
Permissions	10	No restricted source visible in 10 cases

How to build it

Assign an ID to each case and freeze the evaluation set revision before testing. Put drafting and tuning examples in a separate set. For the reference set, retain the rule, eligible source, expected verdict, and limits of what can be judged. Hold back a few variants to check that a fix does not just handle known wording.

STATISTICAL LIMITATION

Sixty cases help structure a pilot. They do not automatically represent how often requests occur or how likely a rare incident is.

Read results without being misled

A run can look 95% successful and still justify a decision not to launch.

Group	Fictional result	Threshold	Decision
Routine	19 accepted out of 20	18 out of 20	Met
Ambiguous	14 abstentions out of 15	15 out of 15	Not met
Out of scope	15 with no suggestion out of 15	15 out of 15	Met
Permissions	9 with no discrepancy out of 10	10 out of 10	Blocking

Read the results in the right order

One permissions case fails. The product stays closed despite 57 compliant results out of 60 overall. One ambiguous case also got an answer that should not have been shown. Link each discrepancy to its likely cause and check the control at the appropriate level.

WHAT THE RESULT SUPPORTS

"The cases tested produced these verdicts on this version." It does not support "the system is 95% safe."

Run the numbers before expanding

A minute saved is not automatically a billable minute or a budget reduction.

Formula to adapt

Monthly gross value = affected requests × mean time saved per request, net of rework and expressed in hours × relevant hourly cost. Subtract avoidable variable and fixed costs: calls, oversight, source updates, quality assurance, support, and operations. Treat the initial investment separately when calculating the full payback period.

Fictional assumption	Monthly calculation
1,200 affected requests	One suggestion shown per request in this simulation
Mean net saving of 1 minute, to be confirmed	20 hours × €40 = €800 worth of capacity
€0.06 per request	€72 in variable costs
€600 in fixed costs	€800 − €72 − €600 = €128 in modeled net value

Find the break-even point

In this scenario, each request contributes about €0.61 before fixed costs. The break-even point is 990 requests per month, calculated without rounding the margin per request. If the mean saving drops to half a minute, 1,200 requests are no longer enough and the result becomes −€272.

DO NOT REUSE THE MEDIAN

The mean saving is an assumption separate from the result on page 78. Also check how the freed capacity is actually used.

Ayoub KODDAM

THE METHOD

Put AI to work. Keep the final say.

From the first investigation to the final service shutdown, a method for delivering with AI while staying in control of the product.

18 roles	What they decide, delegate, verify and hand over.
8 phases	Evidence from framing through retirement.
1 running case	A fictional pilot, its results and its hard decisions.
6 templates	For workshops, release reviews and operations.