# From one workflow to a maintained system.

Every phase ends with evidence, an accountable owner, and a decision. Expansion follows proven usefulness and controlled risk.

## 1. Discover and define

Interview the process owner and users. Map the current workflow, repeated friction, sources, exceptions, and approval points. Agree a baseline, value hypothesis, data owner, and stop conditions.

## 2. Design and assess

Map trust boundaries and roles, choose the deployment model, assess data quality and model licenses, and document assumptions. Agree the retrieval strategy, retention, monitoring, and backup responsibilities.

## 3. Build a bounded pilot

Use approved representative data in an isolated environment. Integrate identity, access filtering, citations, and audit events early. Bound inference cost and time. Keep human approval for consequential actions.

## 4. Evaluate and decide

Use held-out Bulgarian/English questions and domain reviewers. Measure source support, refusal on inaccessible documents, task completion, latency, and cost. Record errors and unresolved risks; decide whether to extend, revise, or stop.

## 5. Deploy and hand over

Validate the target environment, migrations, recovery, account provisioning, logging, and incident contacts. Train administrators and process owners. Release only after agreed acceptance and explicit launch approval.

## 6. Operate and improve

Review drift, data changes, permissions, failed requests, dependencies, and model updates. Re-run evaluation before changing models or retrieval settings. Maintain a change log and review service scope against actual use.

Duration, staffing, cost, and service levels are agreed per engagement. No delivery date or performance outcome is guaranteed by this methodology.