

WORKED EXAMPLES

ATLAS, WORKED END TO END

The completed Canvas and every evidence sheet for Atlas, the running example in the book.

From **Practical AI for Product Owners and Business Analysts** by Madhusudhan Konda · practicalaiforall.com
Free to copy and adapt for your own work.

THE COMPLETED ATLAS CANVAS

A service-desk assistant for routine travel and expense policy questions. The example is deliberately incomplete: permission and source-change tests still have to pass, and value needs pilot evidence.

CANVAS SECTION	COMPLETED ATLAS REFERENCE ANSWER
1. Business problem and outcome	Service-desk colleagues repeatedly locate policy evidence and draft routine answers. Improve handling time and consistency without transferring policy authority.
2. User and workflow	The primary user is a service-desk colleague. Atlas retrieves and drafts; the colleague inspects, edits, sends or escalates.
3. AI task, pattern and ambition	Finding, comparing and grounded drafting through a filtered RAG assistant. Augmentation, not automation. No entitlement decision or write action.
4. Value and success	Capacity release and service improvement measured against handling time, rework, claim support, review effort, escalation, cost and guardrails.
5. Knowledge, data and access	Two approved global policy collections with status, jurisdiction, replacement and permission metadata. Regional synthesis remains excluded.
6. Required behaviour and boundaries	Cited drafts, visible passages, clarification, conflict and no-evidence routes, deterministic prohibited topics and safe service fallback.
7. Human oversight and exceptions	Service desk reviews routine evidence; policy owner resolves conflict and interpretation. Review behaviour and workload are measured.
8. Risk, governance and ownership	Controls cover eligibility, permissions, claim support, stale sources, incidents and change. Named product, policy, engineering, security and service owners.
9. Evaluation and monitoring	Versioned routine, difficult, restricted, prohibited, stale-source and failure cases across component, response, task, workflow and outcome layers.
10. Pilot and delivery roadmap	Controlled 30-60-90 day service-desk pilot. Pre-pilot hard gates, supervised evidence, operating review and an explicit decision route.

Illustrative teaching example from the book, not a deployed customer result. The early sheets (D1 to D3) show Atlas at an earlier stage, before discovery narrowed it to the service desk.

SHORTLISTED-USE-CASE RECORD AND FOUR LENSES SCORE

FIELD	ATLAS
Candidate	An advisory assistant that lets country-office staff find current procurement policy answers with citations - thus reducing repeated questions handled by the policy team.
Evidence so far	Repeated questions, slow responses across time zones and several conflicting policy versions. Question volume and handling time still need a baseline.
Why shortlisted	Clear user pain, a good fit for retrieval and generation, internal use and a realistic human escalation route.
Gate before pilot	Approve the policy corpus, remove conflicting versions, name the knowledge owner and confirm permission rules.
Canvas handoff	Problem, user, desired outcome, policy Q&A pattern, advisory ambition, human owner and knowledge-readiness gap.

ONE-PAGE OPPORTUNITY CASE

FIELD	ATLAS
Problem and user	Country-office staff wait for policy officers to answer recurring procurement-policy questions. Officers repeatedly search several documents, check versions and rewrite similar guidance.
Evidence and baseline	Illustrative two-week sample of 46 questions: 30 routine, 10 needing comparison across sources, 6 needing a policy owner. Routine questions take a median of 12 minutes of officer time; requesters wait a median of nine working hours.
Proposed intervention	Cited answers to routine questions from an approved collection; refuse or escalate when evidence is missing, sources conflict or interpretation is required. No policy rulings, approvals, entitlement decisions or regional exceptions.
Value hypothesis	If Atlas provides cited answers to routine questions at the point of need, median time to a usable answer will fall because staff no longer wait for an officer to repeat an existing answer.
Measures and guardrails	Time to a usable answer, officer handling time, routine questions resolved without drafting; citation correctness, answer support, correction effort, missed escalations and permission violations.
Alternatives considered	Clean and govern the repository; improve enterprise search; an FAQ service; redesign the policy mailbox; keep the current process.
Main assumptions	Enough questions are routine; approved documents contain the evidence; staff try Atlas first; permissions can be enforced; officers can review samples and escalations.
Full cost drivers	Knowledge clean-up and ownership, discovery and engineering, repository and permission integration, evaluation, officer review, adoption, monitoring, support and content maintenance. Expressed as ranges.
Next decision requested	Approve a short evidence and readiness step: confirm the baseline, identify the approved collection and owner, assess quality and permissions, and build a representative evaluation set.

USE CASE FRAME

FIELD	ATLAS
Use case name	Cited guidance for routine procurement-policy questions.
User and trigger	A country-office staff member needs procurement guidance while completing routine work.
Desired outcome	Obtain usable procurement guidance faster without transferring interpretation or approval authority to Atlas.
Current pain	Searching is slow, source authority is unclear and specialists repeat the same retrieval work.
Intervention point	Retrieve approved passages and draft a concise cited explanation.
Approved evidence	Approved procurement policies, eligible country supplements and maintained guidance within the user's access rights.
AI output	Relevant passages, source metadata, a grounded summary and an escalation reason when needed.
Decision owner	The staff member chooses the next step for routine guidance; a policy owner interprets ambiguity or conflict; an authorised approver owns any procurement decision.
Action boundary	Atlas does not approve purchases, select suppliers, change records, edit policy or make binding interpretations.
Human review	Users inspect citations. Specialists receive the question, sources, draft and escalation reason.
Exceptions and prohibited paths	Missing evidence, conflicting sources, exception requests and named high-consequence topics are escalated. Confidential or prohibited requests are refused.
Measures	Time to a usable answer, officer handling time, routine-question coverage, citation correctness, unsupported-claim rate, correction effort, missed escalations and permission violations.

DISCOVERY EVIDENCE PACK

FIELD	ATLAS
Decision sought	Fund document-metadata remediation and a service-desk prototype for two global policy domains, rather than a broad employee assistant.
Riskiest assumptions	Routine questions are answerable from approved sources; regional supplements can be handled; reviewers can check evidence quickly.
Representative set	Routine, ambiguous, exception, conflicting-source, no-evidence and prohibited examples, kept in the users' own wording (see Section 5.5).
Changed use case	Users narrowed to service-desk colleagues; scope narrowed to two global domains; regional coverage sequenced after metadata ownership is resolved.
Recommendation	Proceed with a service-desk prototype for two global policy domains; sequence regional coverage after metadata ownership is resolved.

Representative example set (Section 5.5)

EXAMPLE TYPE	ATLAS EXAMPLE	EXPECTED TREATMENT
Routine	What is the hotel limit for Berlin?	Retrieve current policy and rate table; answer with citations.
Ambiguous	Can I claim a taxi after a late meeting?	Ask for missing circumstances or show the relevant conditional guidance.
Exception	The approved hotel was inaccessible. Can I exceed the limit?	Provide relevant evidence and route for accountable interpretation.
Conflicting sources	The global policy and local supplement state different limits.	Expose the conflict and escalate; do not blend the values.
No approved evidence	What is the rule for an event not covered by policy?	State that no approved answer was found and route the question.
Prohibited or sensitive	Will my disciplinary case affect an expense claim?	Refuse case-specific interpretation and direct the user to the confidential channel.

AI-READY REQUIREMENTS SHEET

FIELD	ATLAS
User, trigger and scope	Service-desk colleagues handling routine global travel and expense questions in two approved domains.
Inputs and context	Original question, user role and access, selected domain, current interaction context and required location when supplied.
Approved sources and access	Published global policy collections designated by the policy owner; exclude draft, expired and unauthorised copies.
Output and evidence	Short guidance, relevant conditions, supporting passages, document title, status and version.
Deterministic controls	Permission filtering, approved-source enforcement, citation presence, prohibited-topic routes and action boundary.
AI quality expectations	Relevant retrieval, supported synthesis, appropriate clarification and non-authoritative language.
Uncertainty and conflict	Ask for missing context; expose conflicting sources; route when approved evidence is absent.
Refusal and escalation	Refuse sensitive case-specific interpretation and send context plus evidence to the controlled channel.
Human review	Service-desk colleagues can inspect, edit, reject and escalate. Policy owner handles interpretation.
Operational and fallback	Show progress, preserve user input, provide retry, search or contact route and measure response time and review abandonment.
Logging, audit and retention	Record question, sources, configuration, output and reviewer action with controlled access and defined retention.
Acceptance approach	Representative routine, ambiguous, conflict, no-evidence, prohibited and failure cases with case and aggregate criteria.

DATA AND KNOWLEDGE READINESS ASSESSMENT

FIELD	ATLAS
Use case dependency	Routine global travel and expense guidance must be supported by approved passages and visible source status.
Approved collections	Published global travel policy, expense policy and named dependent rate tables.
Authority and owners	Global policy owner confirms eligibility; document and platform owners maintain publication and ingestion.
Content quality	Core global material is usable after duplicate removal and contradiction review; regional material remains mixed.
Eligibility metadata	Domain, status, effective date, jurisdiction, authority, owner, access group and source identifier.
Extraction and structure	Headings and prose extract reliably; rate tables require dedicated checks; scanned legacy files are excluded.
Permissions and privacy	Source access groups must filter retrieval and display; personal case data is unnecessary for routine guidance.
Freshness and change	Daily ingestion for policies, tighter validation for rate tables, with removal and failure alerts.
Dependencies	Rate tables and escalation routes must be available; regional appendices are not combined in the first release.
Experiment evidence	Fifty-question source and passage test, difficult-file extraction checks and negative permission tests.
Material gaps and treatments	Enrich metadata and remove duplicates before pilot; control ingestion failures; exclude regional synthesis.
Readiness decision	Proceed to a controlled service-desk pilot for two global domains, subject to the stated remediation and tests.

SOLUTION PATTERN DECISION RECORD

FIELD	ATLAS
User and job	Service-desk colleagues need faster, consistent guidance for routine global travel and expense questions.
Selected pattern	Grounded service-desk assistant with filtered RAG, visible passages and human review.
Why this pattern	Search alone leaves repeated interpretation and drafting; ungrounded generation cannot meet the evidence contract.
Sources and retrieval	Published global collections only, filtered by domain, status, jurisdiction and permission before ranking.
Generation task	Draft a concise, non-authoritative explanation whose substantive claims are supported by cited passages.
Context and memory	Current question, selected domain and explicit clarification within the session; no long-term personal memory.
Tools	No write tools; evaluate one read-only current-rate lookup after dedicated validation.
Action authority	Explain, clarify, refuse and recommend a route; do not approve, update records or make binding determinations.
Human role	Inspect evidence, edit or reject the draft and escalate interpretation or sensitive cases.
Controls and fallback	Deterministic eligibility, citation presence and topic routes; filtered search or human contact on failure.
Observability and evaluation	Trace source versions, retrieval, configuration, output and reviewer action against representative cases.
Excluded scope and next gate	No regional synthesis or agentic action; prove retrieval, groundedness, permission isolation and review effort in pilot.

AI RISK AND CONTROL REGISTER

FIELD	ATLAS
Product boundary	Service-desk guidance for routine global travel and expense questions; no binding interpretation or state-changing action.
Credible failure	Atlas presents an expired or ineligible source as current policy and a colleague sends the guidance.
Affected stakeholders	Service-desk colleague, employee, policy owner and operational support team.
Consequence and blast radius	Incorrect employee action, rework and repeated unofficial interpretation across cases.
Likelihood evidence	Duplicate and ambiguous source statuses were observed; representative retrieval and change tests are required.
Preventive controls	Approved-source catalogue. Status, jurisdiction and permission filtering. Controlled ingestion.
Detective controls	Claim-support evaluation, stale-source alerts, negative permission tests and sampled response review.
Correction and recovery	Reject the response, correct source metadata, re-index, notify impacted reviewers and use fallback search.
Human authority and escalation	Service desk reviews routine evidence; policy owner resolves conflict and interpretation.
Control owners	Product Owner, policy owner, retrieval engineering owner, security owner and service owner.
Evidence of effectiveness	Versioned test cases, access-test results, source-change simulation, incident drill and review-effort data.
Residual-risk decision	Proceed to a bounded service-desk pilot after remediation; regional synthesis and write actions remain excluded.

DELIVERY RESPONSIBILITY VIEW

FIELD	ATLAS
Decision or product behaviour	Only approved current global policy sources may enter Atlas retrieval for the service-desk pilot.
Accountable decision owner	Policy owner for source authority; Product Owner for the pilot boundary.
Required contributors	Knowledge operations, retrieval engineering, security, service desk and risk.
Evidence required	Collection catalogue, status and jurisdiction metadata, replacement mapping, permission tests and retrieval evaluation.
Delivery artefact or system location	Source catalogue, ingestion configuration, AI-Ready Requirements Sheet and evaluation repository.
Approval point	Pilot gate after source eligibility and negative access evidence pass.
Operational owner	Knowledge operations maintains source status; service owner monitors the route.
Incident or escalation owner	Service owner coordinates; policy owner resolves authority or interpretation.
Change authority	Policy owner approves source changes; Product Owner and control owners approve material scope change.
Open dependency	Confirm replacement relationships and complete negative permission tests before the pilot.

EVALUATION PLAN

FIELD	ATLAS
Product behaviour or failure mode	Atlas presents an expired or ineligible source as current policy and a colleague sends the guidance.
Evaluation layer	Component, response, task, workflow and risk outcome.
Representative case set	Routine replacement, stale source, wrong jurisdiction, permission denial and source-change regression cases.
Reference evidence or expected treatment	Current approved source and replacement mapping; stale or ineligible content must not enter retrieval, context or citations.
Measure and method	Eligibility assertions, recall at five, claim-support review, reviewer task observation and source-change simulation.
Threshold or decision rule	Zero ineligible-source inclusion; all high-consequence claims supported; safe no-evidence route when current evidence is absent.
Dataset or case-set owner	Policy owner with the evaluation owner maintaining version and traceability.
Evaluator or reviewer	Retrieval engineering, security tester, calibrated policy reviewer and service-desk participant.
Version and configuration under test	Recorded source, index, retrieval, prompt, model, validator, interface and evaluation-set versions.
Result and failure pattern	Pending pre-pilot run; failures grouped by eligibility, retrieval, generation, review and change.
Monitoring signal	Stale-source alert, sampled claim support, permission-denial pattern and reviewer correction.
Gate, corrective action or stop decision	Proceed only after hard gates pass; otherwise sequence remediation, reshape the route or stop.

30-60-90 DAY VALIDATION ROADMAP

HORIZON	SCOPE, CAPABILITY AND EVIDENCE	OWNERS, GATE AND DEFERRAL
Days 1 to 30	Stabilise two global collections, pilot environment and filtered RAG. Complete permission, stale-source, prohibited-route, offline evaluation and incident drills.	Product, policy, engineering, security and evaluation owners. Gate: hard controls and high-severity cases pass. Defer regional, employee and write capability.
Days 31 to 60	Supervised service-desk use. Measure task success, claim support, evidence inspection, review effort, escalation, latency, cost and support demand.	Service owner, design, change, policy and evaluation owners. Gate: useful workflow without unacceptable review or control failure.
Days 61 to 90	Bounded pilot operation, source-change exercise, recovery test and business-risk review. Compare the baseline and prepare scale decisions.	The accountable gate owner reviews evidence. Decide using the five decision routes. Any expansion is separately bounded.