

AI Solution Strategy

A content engine for generating and updating microlearning materials for precinct election commission members and citizen observers

Fundacja Odpowiedzialna Polityka (Political Accountability Foundation, Poland)

„Odporna Demokracja” / Resilient Democracy programme

Document version	2.0
Date	20 August 2026
Content owner	Subject-matter lead, Fundacja Odpowiedzialna Polityka
Project manager	Dominika (FOP)
Programme	AI Impact Lab · Tech To The Rescue
Legal sources	National Electoral Commission Guidelines (Resolution No. 211/2023, as of 25.09.2023) · Electoral Code (as of 31.03.2023)
Working solution	materialy-edukacyjne-v1.vercel.app

Every claim in this document is traceable to a source listed in Annex A, or is flagged as an open question with a named owner and a next step (Annex B). Responsibility for the content and for all decisions remains with the Foundation.

A note on terminology

Polish electoral roles do not map one-to-one onto English usage. This document follows the terminology used in international election observation:

Polish term	English used here	Meaning
Obwodowa komisja wyborcza (ObKW)	Precinct Election Commission (PEC)	The body that runs voting and counts ballots in a single polling station
Członek / członkini komisji	Commission member	A citizen appointed to serve on a PEC, usually for a single election
Obserwator społeczny	Citizen observer	A non-partisan observer delegated by an eligible civic organisation
Mąż zaufania	Party agent	An observer delegated by an election committee, with a different and broader set of formal rights
Państwowa Komisja Wyborcza (PKW)	National Electoral Commission (NEC)	The central election management body; issues binding guidelines before each election
Kodeks wyborczy	Electoral Code	The statute governing elections in Poland

The distinction between a **citizen observer** and a **party agent** is not cosmetic. A citizen observer is impartial, does not represent any committee, does not intervene in the commission's work and cannot enter remarks into the official protocol. That boundary is anchored in Article 103c § 2 of the Electoral Code and is treated in this project as non-negotiable content.

Table of contents

Executive summary

1. Project basics

2. Problem description and scope

2.1 The problem · 2.2 Scope · 2.3 Out of scope · 2.4 Deliverables

The decisive choice — what counts as a unit of review

3. AI learning goals and rationale

4. Success metrics and social impact

4.1 Five metrics · 4.2 Reach · 4.3 A precondition for credible measurement

5. Assumptions, constraints and dependencies

6. User flow

6.1 The FOP reviewer · 6.2 The learner

7. Maintenance plan

8. Requirements (MoSCoW)

9. Ethical considerations and risks

9.1 AI-related risks · 9.2 Organisational risks

10. Resources and planning

11. Governance and implementation

11.1 Roles and decisions · 11.2 Legal considerations · 11.3 Budget · 11.4 Implementation plan

Annex A — evidence register

Annex B — open questions

Annex C — handover package

Executive summary

Fundacja Odpowiedzialna Polityka trains the two groups on whom the integrity of Polish elections depends: members of precinct election commissions and citizen observers. Their training materials must be fully faithful to the binding guidelines of the National Electoral Commission and to the Electoral Code — and those change, usually a few weeks before election day, which is the Foundation's busiest moment.

We built an engine that does this work in a repeatable and verifiable way. It drafts content strictly from three approved inputs, routes every item through a mandatory expert review, and publishes only what a human has approved — with a visible version, effective date, owner and reference to the provision the content rests on. Where no basis exists in the approved sources, the engine stops and flags the gap instead of guessing.

Where the project stands. The prototype works and is publicly available. An external expert reviewed and published 21 items across two content paths, each with full provenance. The engine covers eight stages of election day and two roles. The generated corpus comprises 320 quiz questions, 52 checklist items, 44 microlearning questions with complete legal anchoring, and 15 interactive scenarios.

What the Foundation's own data shows. Materials for commission members are the Foundation's most heavily downloaded resource: in the 2023 parliamentary elections the commission-members section drew 27.4% of all traffic, and in 2025 it accounted for 56.5% of all key events while representing only 17% of page views. Mobile traffic rose from 31% to 57%. Practically the entire value of a material is realised within a window of roughly 72 hours around election day.

One decision ahead. Generating content is cheap. Two things are not: anchoring every statement in a specific legal provision, and a human signature of accountability. Whether the unit of review remains a single content item, or becomes an *anchored pattern* built on a point of law, determines whether the cost grows without limit or is capped by the number of source provisions — and therefore whether the materials will be ready before the 2027 elections. The evidence that the second approach works already exists: questions drafted from a guideline point carry a complete set of legal references, questions drafted from the content itself carry none.

1. Project basics

Field	Value
Name	Content engine for generating and updating microlearning materials — checklists, quiz questions and interactive scenarios for PEC members and citizen observers
Content owner (signs off publication)	Subject-matter lead, FOP

Field	Value
Project manager	Dominika (FOP)
Build, deployment and testing	Carlos (AI Impact Lab)
Content reviewer	Expert 5 — reviewed and published 21 items (24.07.2026)
Outreach to end users	FOP communications team (final stage)
Content generation support	Claude (AI Impact Lab)
Live solution	https://materialy-edukacyjne-v1.vercel.app/
Reproducible artefact	A self-contained <code>index.html</code> file from the main branch of a public repository — runs in a browser with no installation, no backend and no login
Timeline	Clickable prototype: delivered (August 2026). Pilot: date to be set by the FOP content owner

2. Problem description and scope

2.1 The problem

Training materials for commission members and citizen observers must keep pace with changes to the NEC guidelines and the Electoral Code, and must remain fully faithful to the legal source. Today updating them is manual, slow and error-prone. There is no repeatable process that shows, for every piece of information, where it comes from: which provision, in which version, and who approved it.

A second layer of the problem concerns **format**. Five existing materials described as checklists run to 8–13 pages, with an average sentence length of 30–37 words; they do not function as checklists. Mobile traffic rose from 30.9% in 2023 to 56.9% in 2025, overtaking desktop for the first time. The entire value of the materials is realised in a window of about 72 hours: on the Saturday before the 2023 election there were 2,104 downloads, 1,306 on election Sunday, and 121 on the following Monday — a drop of 94%.

The problem is therefore not only the speed of updating, but whether the finished material can actually be used in a polling station, on a phone, within those 72 hours.

2.2 Scope

We are building the content back office and the interface that operates it. An FOP reviewer selects the audience, the stage of election day and the type of material; the engine drafts a proposal from three inputs (learning objectives, the Foundation’s own materials, legal

sources); the reviewer approves, requests a correction, or rejects each item individually; only approved items enter the database — with version, effective date and owner.

2.3 Out of scope

- A finished learning application with user accounts, or an adaptive training platform.
- Automatic publication without human involvement.
- Interpretation of law by AI, resolution of legal doubts, and any citation not grounded in an approved source.
- The English-language corpus as a product for end users (see Annex B).

2.4 Deliverables

State verified in the production file on 20.08.2026:

Deliverable	Status	Count	Legal anchoring
Quiz question bank, Polish	generated, awaiting review	320 (2 roles × 8 stages)	0 / 320 — every item carries the placeholder [TO BE COMPLETED: point/article]
Quiz question bank, English	parallel corpus	320	0 / 320 — same placeholder
Microlearning questions, stages A–H	generated from guideline points	44	44 / 44 — with the specific provision and its legal status date
Checklists	generated	52 items	52 / 52
Interactive scenarios	3 carried over from FOP courses, 12 generated	15	marked as unverified — awaiting review
Items published by the reviewer	signed off by Expert 5	21	21 / 21 with source, version, date and reviewer name

The placeholder is not a defect — it is a safety mechanism. It blocks publication of any item for which the AI found no basis in a permitted source, instead of allowing it to guess a provision number.

The decisive choice — what counts as a unit of review

This single decision drives the budget, the timeline and whether the materials will be ready in time. This document does not settle it; the FOP content owner does. The evidence is set out below.

The real cost axis

Cost here does not split along the usual line between manual work and automation. It splits like this:

Does cost scale with the number of **content items** (an unbounded number), or with the number of **points of law** (a bounded number — 195 guideline points, realistically 40–60 of them concerning election day)?

Generating content is cheap: in the prototype it happens with no model calls at all, on a pre-built corpus. Two things do not scale: anchoring a statement in a provision, and the human signature of accountability.

Two options

Dimension	A. The unit is a content item	B. The unit is an anchored pattern (point of law + question type + typology of misconceptions)
What the expert approves	every question and task individually	the pattern, a sample check, and the version
Marginal cost of a new item	high and constant	close to zero
Current corpus of 320 items	roughly 15 expert sessions	40–60 points to map, then generation
A second language version	doubles the full review cost	doubles only the pattern layer
An amendment to the guidelines	cost comparable to building from scratch	bounded by the number of amended points
Precondition	none — it works today	anchoring must first reach 100% at the level of the point of law
Risk	elections have a fixed date; supply of content will not meet demand	a flawed pattern propagates across many items — sample checks are required

Evidence from the production file

Two corpora sit side by side in the same file, produced by two different methods:

- **320 quiz questions** drafted “from the question” — **0% carry a specific legal provision**;
- **44 microlearning questions** drafted “from the guideline point”, working from fragments of the legal sources — **100% carry the provision and its legal status date**.

This is not a difference in model quality. It is a difference in the direction of work. Option B is not a hypothesis: it has already been executed once, on 44 items.

Recommendation and the missing data point

Editorial recommendation: **option B as the target direction, option A as the fallback for contested items**. The knowledge needed to build the patterns already exists in the prototype — 804 explanations of typical misconceptions, currently scattered across items instead of being collected by point of law. The work is a change of ordering, not creation from scratch.

To make the decision calculable, one data point is missing: **the time and type of reviewer decisions from a single full session** — minutes per item, broken down into “approved”, “substantive correction”, “editorial correction” and “rejected”. If more than 70% of corrections turn out to be substantive, patterns are the only way forward; if they are editorial, the problem lies in how the AI is instructed. Cost of obtaining this data: one change to the review sheet. Measurable at the next session.

3. AI learning goals and rationale

What the Foundation is learning. How to constrain a model to an approved register of sources; how to phrase instructions so that content can be corrected after review; and how to combine AI drafting with mandatory human control and an audit trail. This knowledge is transferable — the Foundation can share it with other election-focused organisations and with other non-profits adopting AI responsibly.

Why AI is the right tool. Language models handle repeatable generation of structured content from several inputs at once, and rapid updating when a source changes — provided they stay strictly within approved material. The factor that makes the solution scalable is precisely this controlled, resumable work, not autonomous legal interpretation: one change in the law triggers the chain of draft → review → publication with a stated version.

What the AI does not do here. It does not resolve legal doubts, does not classify conduct, and does not guess the number of a provision. The engine’s behaviour on hard inputs has been documented in two cases: refusal to generate outside the approved scope, with the topic logged to a backlog; and, where no basis exists in the sources, flagging the gap, marking the item “pending legal verification” and referring it to the expert.

4. Success metrics and social impact

4.1 Five metrics with a baseline

Metric	Today	How and when measured	Target
Completeness of legal references	quizzes 0/320 · microlearning 44/44 · checklists 52/52	database audit, before and after review	100% before publication — already enforced by the publication block
Coverage: role × stage × material type	21 items published, 2 combinations	coverage metric in the database, monthly	growth and parity between the two roles; level set by the FOP owner
Review throughput	21 items per session; time per item not measured	timestamp in the review sheet	to be set after the first measured session
Structure of reviewer decisions	unknown — reasons for corrections not recorded	decision log with justifications	substantive corrections below 30%
Reach within the 72-hour window	2023: 2,104 downloads on Saturday, 1,306 on Sunday, 121 on Monday	analytics tool, per election	conditional on the change described in 4.3

4.2 Reach

The Foundation's own data for 2023–2025 shows that the group this engine produces content for is its most strongly engaged audience:

- **2023 parliamentary elections, 15-day window:** 34,810 page views, 13,091 users, 138,463 events.
- **“For commission members” section:** 9,547 views — 27.4% of all traffic — and 4,288 users; the most visited section of that election.
- **2025 presidential election:** the commission section accounted for 3,362 of 5,949 key events (56.5%) while drawing only 17.0% of views. Together with the commission handbook and the observer section, this represents 85.1% of all conversions.
- **Devices:** mobile share 30.9% (2023) → 33.4% (2024) → 56.9% (2025).

A projection of the number of beneficiaries after deployment remains open (Annex B); the most reliable basis will be the download count of the Foundation’s “Essentials” guide.

4.3 A precondition for credible measurement

Event-data retention in the Foundation’s analytics tool is set to two months, which means download figures for individual materials can no longer be recovered for any past election. Changing the setting to 14 months is a single click at zero cost — but it works only going forward. Without it, the reach metric will remain unverifiable after the 2027 elections as well.

Social impact is therefore measured on two tracks: the quality and currency of the content (metrics 1–4, own data, high reliability) and reach (metric 5, reliability conditional on the change above). The FOP content owner is accountable for the data.

5. Assumptions, constraints and dependencies

Type	Description
Assumption	An FOP reviewer is always in the loop. Unapproved content is an engine proposal, not a training material
Assumption	The expert is available in the window after the NEC guidelines are published — that is, when the Foundation is busiest
Constraint	The engine works only on the approved register of sources; outside it, it stops rather than guesses
Constraint	Work in progress is not persisted within a session — it must be exported to a file
Constraint	The only non-scalable resource is the subject-matter expert’s time: 320 items amount to roughly 15 sessions at current throughput
Dependency	Currency and versioning of the legal sources; the “what has changed” module requires a new version of the guidelines to be loaded
Dependency	A second infrastructure administrator — position currently unfilled

Type	Description
Scaling barrier	Without a measurement of review time per item there is no cost model, and therefore no basis for a funding application

6. User flow

6.1 The FOP reviewer — production path

1. Select the audience (commission member or citizen observer), the stage of election day (A–H) and the material type (checklist, questions, interactive scenario).
2. The engine drafts a proposal from three inputs: learning objectives, FOP’s own materials, legal sources.
3. Review item by item: **approve** / **request correction** / **reject**; the latter two require a justification.
4. Only approved items are published — with version, effective date and owner; the previous version moves to the archive and the user is informed what has changed.
5. Materials database and trial mode, with review status visible and a dual provenance label: material from an FOP course, engine-generated content, or an item from the approved bank.

Safe path: where no basis exists in an approved source, the engine halts generation and the topic goes to a backlog; the decision to extend scope belongs to a human.

6.2 The learner

A learning unit takes 10–12 minutes and consists of a scenario, five questions and a drag-and-drop ordering of a checklist (with keyboard operation for users who do not use a mouse). The scenario inherits its legal content from previously approved items; the AI is responsible only for the fictional framing, never for the provision.

This layer answers a gap identified by users themselves in the Foundation’s research: the absence of a “what to say and what to do” layer — concrete scripts of conduct for difficult situations in a polling station.

The boundary of the citizen observer’s role is inviolable within it: **recognise** → **record factually** → **report through the proper channel**. The observer does not intervene in the commission’s work, does not resolve doubts and does not enter remarks into the protocol (Article 103c § 2 of the Electoral Code).

7. Maintenance plan

The FOP content owner is accountable for maintenance. The trigger is the publication or amendment of NEC guidelines or of the Electoral Code: the engine proposes updates to the

affected items, the reviewer checks them, publication receives a new version and effective date, and the previous version remains in the archive.

Two things must be completed before the pilot: the procedure for loading a new version of a source, and the full “what has changed” module — a concise summary of the differences between versions plus a separate quiz on the changes. At present only a change notification is implemented.

A separate, open item is the exception procedure for corrections issued by the NEC or the Supreme Court after the training content has been frozen.

8. Requirements (MoSCoW)

Must have

1. Generation strictly from the three approved inputs; no interpretation of law and no guessing of a provision number.
2. Mandatory item-by-item review, with a justification required for corrections and rejections.
3. A hard publication block on the gap placeholder — in both the Polish and the English corpus.
4. Version, effective date and owner on every item; archiving of the previous version; an audit trail.
5. A dual label on every element: content provenance and review status.
6. Scenarios inherit their legal content from approved items; the AI is responsible only for the fictional framing.
7. Generation halts outside the approved scope and the topic is logged to a backlog.
8. The 15 scenarios marked as unverified do not reach learners before review — they look finished, and that is precisely the risk.

Should have before the pilot

- Measurement of the time and type of reviewer decisions — a precondition for the cost model and for the unit-of-review decision.
- A recorded, complete run of a rejection, including the justification.
- Legal anchoring of the quiz corpus brought to completion under the chosen option.
- Persistent storage of work in progress — this unlocks parallel work by several reviewers.
- A second infrastructure administrator appointed, with a handover instruction.
- A one-page checklist card designed mobile-first, as a distinct material type.
- A decision on the status of the English version.

Could have

- The full “what has changed” module: a summary of differences between versions plus a dedicated quiz.

- A learner-facing panel.
- Export of approved units to the Foundation's course platform.

Will not have in this phase

- A learning application with user accounts. That is the step from artefact to service: personal data, accessibility, user support on election day and peak load within the 72-hour window. This boundary is a deliberate decision.
- Publication without human involvement.
- The English version as a product for end users — translating a statement of law is a legal act, not an editorial one.

9. Ethical considerations and risks

9.1 AI-related risks

Risk	Likelihood	Mitigation
An incorrect reference to a provision	medium	No guessing; gaps are flagged; publication is blocked; source and legal status date on every item; review before publication
The block does not cover the placeholder in the English corpus	medium	Verify and close before the pilot
Content looks finished although it has not been reviewed	medium to high	Visible status on every item, publication block, review queue for scenarios
Confusing the citizen observer role with that of a party agent	medium	Audience labelled on every item; the role boundary anchored in Article 103c § 2; scenarios teach the reporting path, not intervention
Incorrect content reaching users unchecked	low	No approval decision means no item in the database
A flawed pattern propagating across many items	medium, conditional	Sample checks and the owner's signature on the version

9.2 Organisational risks

Risk	Likelihood	Mitigation
No measurement of review time, hence no cost model	high	Timestamp at the next review session
No recorded evidence of the rejection path	medium	Execute and archive before the pilot
Knowledge and access concentrated in one person	medium	Appoint a second administrator; prepare a reproduction guide
Material format unsuited to the realities of election day	medium	One-page card designed from the phone screen
Accessibility for people with disabilities in the learner-facing layer	high	Every interactive package examined at the Foundation lacks heading structure and a declared document language, and permits zoom to be disabled on mobile; the new layer must not inherit these defects
Outdated training materials still reachable outside the prototype	medium	Legal review of older courses and a decision on withdrawing them
Retention of personal data on the course platform	medium	To be resolved by the Data Protection Officer

10. Resources and planning

Tools and environment. A single self-contained HTML file holding all the logic, styled in the Foundation's visual identity. It runs in a browser with no backend and no login. Published on Vercel (free tier), with the code in a public repository; changes are deployed via a working branch, preview and merge. Data can be exported from the browser to CSV, XLSX, HTML and JSON. There are no integrations with external systems. In the prototype, generation is replayed from a pre-built corpus without model calls.

Data sources. NEC Guidelines / Resolution 211/2023 (as of 25.09.2023) and the Electoral Code (as of 31.03.2023), held as fragments: 195 and 611 units respectively. The Foundation's own materials: the checklist template, the list of principal commission tasks,

the Citizen Observer Code of Conduct (10 rules), the “Do pytań” question series, the 2026 learning objectives, and scenarios from FOP courses. Pilot data: the 15 October 2023 parliamentary elections and referendum. All sources are either public or owned by the Foundation.

Data sensitivity. Training materials contain no personal data. The prototype does not log users and does not persist session state, so data protection law does not apply to the materials themselves. That changes on the day a learner-facing panel with accounts is launched — which is why the panel is deliberately outside the scope of the present phase.

Data gaps. The legal provision in the 320 quiz questions remains to be completed by the reviewer — deliberately, without guessing. No amended version of the guidelines has been loaded, so the “what has changed” module cannot yet be demonstrated on a real difference. There is no measurement of the time and type of reviewer decisions. Fifteen scenarios await source verification. Historical download data for individual files is irrecoverable.

Test plan. The second-person test has been carried out and accepted: Expert 5 reviewed and published 21 items across two paths with full provenance (24.07.2026). The criteria were fidelity to the source, soundness of the review decisions and completeness of metadata. Technical testing at the most recent deployment: 75 of 75 local assertions and 17 of 17 production checks. The next test will cover a complete run of a rejection, a review of scenario items against the boundaries of the observer role, and a review session with time measurement.

11. Governance and implementation

11.1 Roles and decisions

The project owner, who signs off publication, is the FOP subject-matter lead. Day-to-day management rests with Dominika. Carlos is responsible for build, deployment and testing. Content review is carried out by Expert 5. The team works part-time on the project; for the Foundation’s leadership it is a priority.

11.2 Legal considerations

Content must be faithful to the electoral law in force, with the legal status date visible on every item. The AI does not interpret provisions. Training materials contain no personal data. The solution operates as an assistive system with mandatory human control and no autonomous AI decisions, which corresponds to the requirements of EU artificial intelligence legislation for this class of application. The compliance assessment must be repeated when a learner panel with user accounts is launched.

11.3 Budget

The figures below are the Foundation’s own estimate and require approval. Prices for team access to AI tools should be verified against the provider’s current published pricing. Euro equivalents should be added at the exchange rate applicable on the date of submission.

Fixed costs

Item	Cycle	Amount (PLN)	Notes
Prototype hosting	one-off	0	HTML file on a free tier
Team access to an AI tool (min. 5 users)	annual	7,200	to be verified against current pricing
Coordinator	monthly	2,000	assuming 20 hours per month
Technical expert	monthly	2,000	assuming 20 hours per month
Subject-matter expert (review)	lump sum	8,000	content sign-off after the guidelines are published
Outreach, 6 months to election day	total	6,000	staff time and paid promotion
Visual concept for the learner-facing version	one-off	1,000	prepared jointly with AI

Variable costs

Driver	Unit	Cost	Notes
Expert review time	per item	not measured	The dominant variable cost of the whole undertaking. Throughput is known (21 items per session); the per-item time is not
AI model usage	per generation run	marginal	Several iterations of the corpus cost a fraction of one hour of expert time
Production hosting	monthly	approx. 80	if the free tier ceases to suffice

Stage forecast (estimate, subject to approval): prototype, 0–3 months, approx. PLN 20,200 · pilot, 3–6 months, approx. PLN 12,000 · scaling, 6–12 months, approx. PLN 40,000. The funding source remains open.

Until the review time per item is known, the scaling forecast is an estimate rather than a model. That single data point changes all three figures.

11.4 Implementation plan

Stage	Timing	Milestones
Prototype	delivered	Clickable prototype with the full working path, safe stop and review gate; 21 items published by an external expert; microlearning for all eight stages; 15 interactive scenarios
Pilot	3–6 months	Review-time measurement · recorded evidence of the rejection path · persistent storage of work · parity of coverage across both roles · a controlled live test from the Foundation's own accounts
Scaling	6–12 months	The unit-of-review decision implemented · full “what has changed” module · one-page mobile card · the solution's framework handed to other election-focused organisations

Annex A — evidence register

#	Evidence	Source	Date
A1	21 items published with full provenance	CSV exports from the	24.07.2026

#	Evidence (4 checklist items for commission members, 17 questions for citizen observers), version 1.0, reviewer Expert 5	Source materials database	Date
A2	Inventory of 94 prototype features: 78 working, 5 to be built, 1 deliberately out of scope; throughput of 21 items per session; 52 of 52 checklist items anchored; zero model calls	prototype cost analysis	04.08.2026
A3	Data on traffic, devices, the 72-hour rhythm, user ratings of materials, accessibility and portfolio gaps	audit of FOP educational materials, with inventory	07.08.2026
A4	Production deployment: 75 of 75 local tests, 17 of 17 production checks	deployment record	18.08.2026
A5	State of the production file: 320 Polish questions (0 with a provision), 320 English, 44	read of the file from the repository's main branch	20.08.2026

#	Evidence	Source	Date
	microlearning questions (44 with a provision), 15 scenarios, 563 gap placeholders		
A6	Documented engine behaviour on hard input: refusal outside scope, and gap flagging with referral to the expert	hard-case documentation	—

Annex B — open questions

#	Question	Owner	Next step
1	Unit of review: content item or anchored pattern	FOP content owner	decision after the first measured review session
2	Does the publication block cover the gap placeholder in the English corpus	Carlos	test on an English item before the pilot
3	Is the English version a product for users or a demonstration layer	FOP content owner	decision; it determines whether review cost doubles
4	Formal confirmation of the division of responsibility: who signs off publication, who manages the project	Dominika and the FOP content owner	one sentence for the document's front matter

#	Question	Owner	Next step
5	Pilot date and funding source	FOP content owner	enter the date
6	Exception procedure for NEC or Supreme Court corrections issued after content freeze	FOP content owner	separate governance thread
7	Change analytics event retention from 2 to 14 months	FOP analytics administrator	a single setting; effective going forward only
8	Projected number of beneficiaries after deployment	FOP content owner	estimate based on downloads of the “Essentials” guide

Annex C — handover package

What the solution does. It drafts microlearning content from three approved inputs, routes it through item-by-item review, and publishes only what a human has approved — with version, effective date, owner and visible provenance.

What it does not do. It does not train end users, does not publish without a human, does not interpret law and does not guess provision numbers.

Location. Live solution: <https://materialy-edukacyjne-v1.vercel.app/>. Reproducible artefact: the `index.html` file from the main branch of the public repository. Review evidence: CSV exports of 24.07.2026. Accompanying documents: this strategy, the prototype cost analysis, and the audit of educational materials.

Reproduction instructions. Download `index.html` and open it in a browser — no installation, backend, login or network connection required. To deploy a change: working branch, preview, merge to main; publication is automatic.

Owner of further work. The FOP content owner (content and publication) and Carlos (repository and deployment) — until a second administrator is appointed.

Next test. A review session with per-item time measurement, and a complete run of a rejection.

Escalation. A legal doubt goes to the FOP expert and the item stays blocked until it is resolved. A question of scope goes to the project owner through the backlog. A failed deployment means reverting to the previous version.

Kill switch. Roll the publication back to the archived version in the database and, in the last resort, withdraw the deployment and work from the local file. None of these operations requires access to external systems.

This document contains no personal data and no confidential information. All figures derive from the artefacts listed in Annex A. Responsibility for the content and for all decisions remains with Fundacja Odpowiedzialna Polityka.