Leitspur

SME AI SYSTEMS

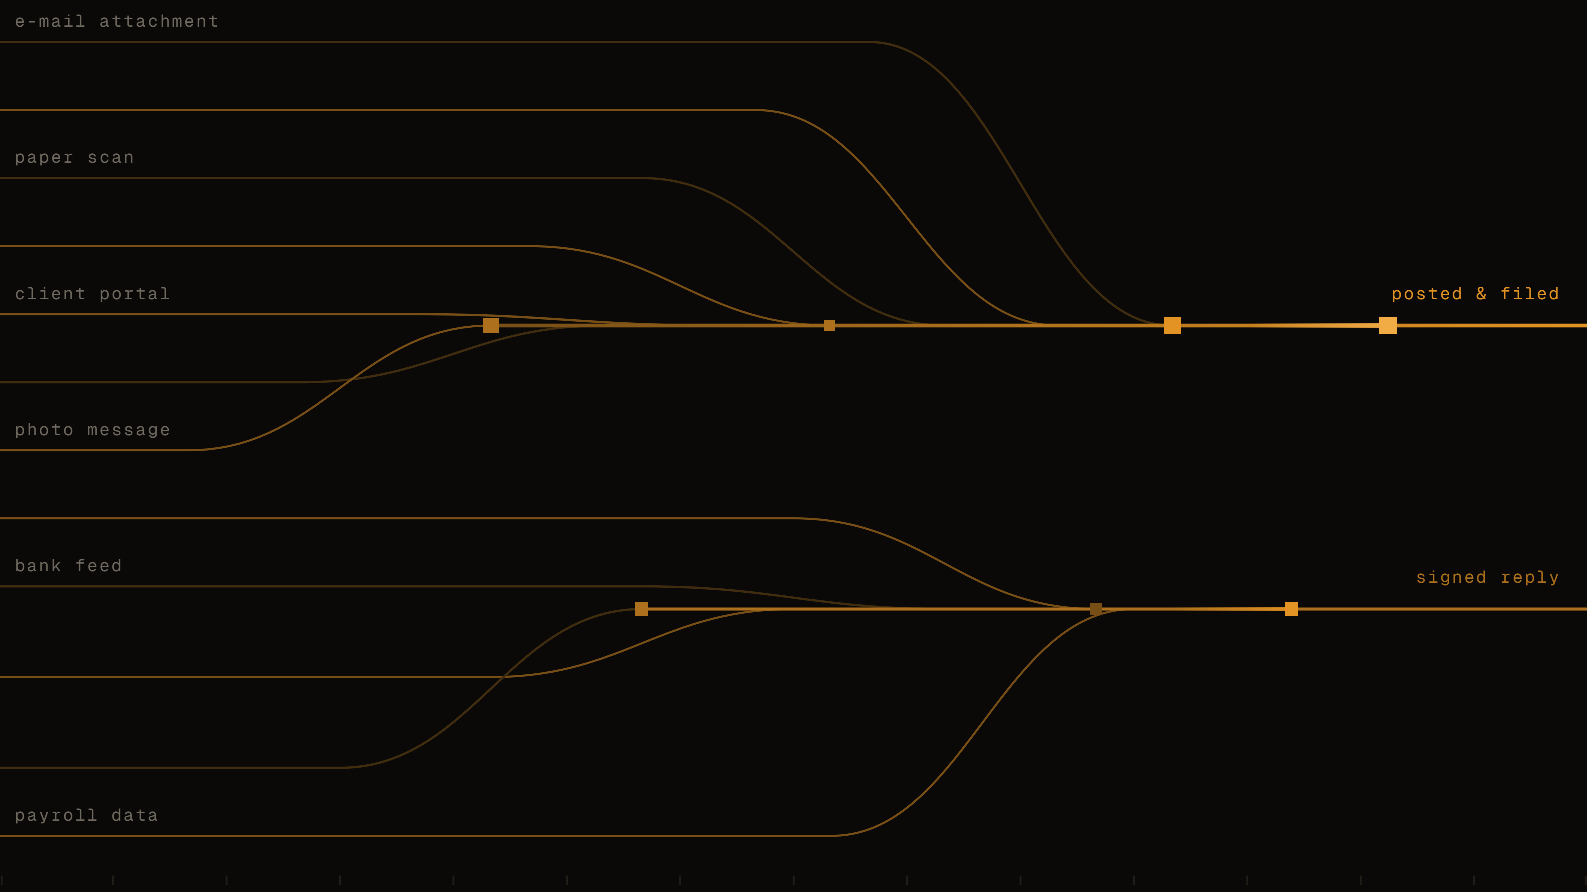

AI READINESS & WORKFLOW MAPPING · SAMPLE REPORT

The mandate workflow and its readiness for AI.

PREPARED FOR

Alpstein Treuhand AG

St. Gallen · July 2026

A fictional company. Scope, method, and depth match a real engagement.

REFERENCE	FIELDWORK	ISSUED	STATUS
LS-SAMPLE-02	July 6 – 17, 2026	July 27, 2026 · v1.0	Sample report — free to share

ABOUT THIS ENGAGEMENT

Scope, method, and contents of this report.

Alpstein Treuhand AG asked Leitspur to map the mandate workflow from client document intake through bookkeeping, filings, and correspondence, assess AI readiness, and design a target workflow that reduces mechanical load without affecting the judgment work clients actually pay for.

Method

- Nine interviews across both partners, fiduciary specialists, accountants, and payroll — deliberately including the most skeptical voices.
- Three half-days of observation during a normal July week: intake sorting, posting runs, VAT reconciliation, client e-mail.
- Systems walkthrough: Abacus and AbaWeb, Bexio (smaller mandates), Dr. Tax, DeepBox archive, the shared Outlook inboxes, and the deadline Excel.
- Quantitative baseline from document counts (April–June 2026), posting statistics, and a categorized sample of 250 inbound client e-mails.

Scope

In scope: document intake, bookkeeping and VAT, client correspondence, deadline management. **Out of scope by agreement:** tax optimization advice itself, audit-adjacent work, and the pension fund administration mandate, which is noted where relevant but not assessed. The plan is written to be executable by internal IT, any competent integrator, or Leitspur; its requirements and gates apply regardless of who builds.

> HOW FIGURES WERE DERIVED
>
> Volumes come from your systems and archives; effort figures were logged by the team itself during one measured week in June, then sanity-checked during our observation days. Franc figures use a loaded internal rate of CHF 105/h agreed with the partners.

Contents

Company snapshot

Company	Alpstein Treuhand AG, St. Gallen — fiduciary services, founded 1994, partner-owned
People	19 employees: 2 partners, 7 fiduciary specialists, 6 accountants, 2 payroll, 2 apprentices
Mandates	≈ 410 active: bookkeeping ≈ 240, payroll ≈ 95, tax returns ≈ 180 (overlapping)
Systems	Abacus + AbaWeb, Bexio (≈ 60 small mandates), Dr. Tax, DeepBox, shared Outlook inboxes
Seasonality	January–April closing season; overtime peaks 28% above baseline

EXECUTIVE SUMMARY

The team spends roughly a third of its working time sorting and typing.

3,650 docs/mo CLIENT DOCUMENTS RECEIVED, Ø APR–JUN 2026	**58** % ARRIVE UNSTRUCTURED: E-MAIL PROSE, PHOTOS, PAPER	**310** h/mo SORTING AND KEYING BEFORE ANY ACCOUNTING BEGINS	**+28** % OVERTIME PEAK IN CLOSING SEASON (Q1)

What we found

- Five intake channels — e-mail, paper, portal, photo messages, bank feeds — all converge on **one manual sorting desk**. It absorbs 310 hours a month, much of it worked by trained accountants.
- **71% of postings follow recurring patterns** the team recognizes on sight (same supplier, same account, same VAT code), yet each one is keyed by hand every month.
- Of 840 inbound client e-mails a month, our 250-e-mail sample shows **roughly 60% ask recurring questions** (VAT treatment, payroll certificates, deadline confirmations) whose answers exist in the mandate file.
- Deadline tracking depends on **one Excel file and two people**. 2025 saw 14 documented near-misses; none led to fines, but each caused overtime.
- Practice varies by specialist: the same document type is filed three different ways. This is the main obstacle to automation, and it also slows the onboarding of new staff.

What we recommend

- Build an **AI intake and posting line**: every document classified, extracted, and routed by confidence. Recurring cases are auto-drafted for batch approval; unusual cases go to a specialist, as they do today.
- Start with two **trust-building quick wins**: drafted client replies (specialist signs every one) and a deadline radar that replaces the Excel.
- Standardize filing practice *as part of* Phase 1; the automation requires the consistent filing the partners have wanted for years.

READINESS VERDICT

2.9 / 5. Systems and data discipline are ahead of most firms this size. The open questions are organizational: two senior specialists are openly skeptical, and confidentiality must be settled in writing before any client document reaches a model. The plan addresses both points.

Expected effect at full rollout

METRIC	TODAY	TARGET
Sorting & keying effort	310 h/mo	≈ 120 h/mo
Recurring postings keyed by hand	≈ 100 %	≈ 25 %
Ø response time, client queries	1.8 d	same day
Capacity released, team-wide	—	≈ 280 h/mo

≈ 280 h/mo is roughly the working time of 1.7 full-time staff, in a hiring market where fiduciary specialists are hard to find. It is capacity for about 40 additional mandates, or for advisory work that is currently deferred. Conservative year-one value: ≈ CHF 235,000 (p. 10).

HEARD ON SITE

"January to April we live in the scanning room. That is not why anyone joins a Treuhand." — fiduciary specialist

CURRENT STATE

Five intake channels converge on one manual sorting desk.

Everything a client sends, in whatever form, passes through manual sorting before any accounting can begin. The funnel below is drawn from the April–June document counts and one measured week of team time logging.

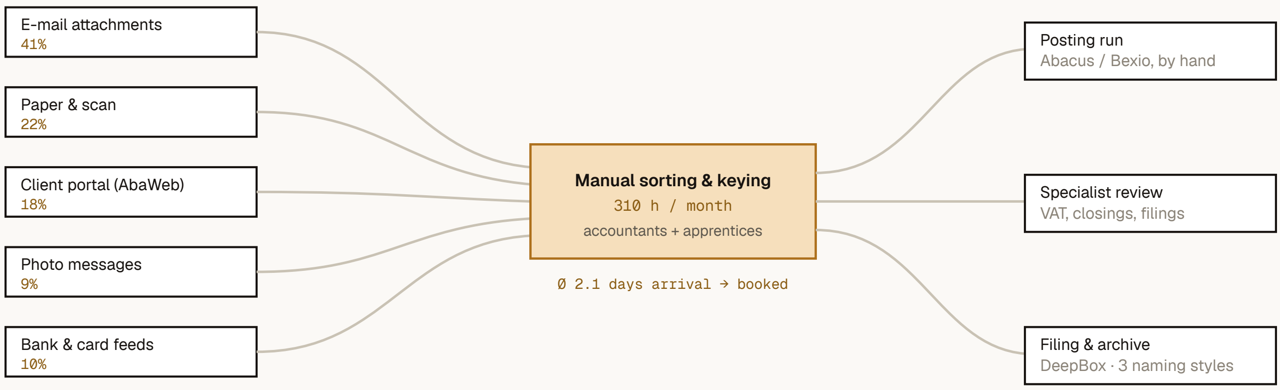

DOCUMENT FLOW, ALL MANDATES. SHARE LABELS = DOCUMENT VOLUME BY CHANNEL, APR–JUN 2026 (N = 10,940).

Observations that shape the design

- **The funnel is also the bottleneck for clients.** A document photographed on Sunday is booked on Wednesday. Clients notice the lag most at VAT deadlines.
- **Recurring patterns dominate.** 71% of the ~28,500 monthly postings match a pattern the team recognizes instantly. This is the standard case for proposal-plus-approval automation.
- **Correspondence is a hidden workload.** 840 e-mails a month at Ø 16 minutes is 224 hours, none of it planned or visible in any system.
- **Filing practice varies by person.** Three naming conventions coexist in DeepBox. This causes staff no difficulty but prevents reliable automation; standardization is therefore Phase 1 work.

QUANTIFIED BASELINE

Rule-based work accounts for 880 hours per month.

Mechanical effort per month by activity — hours

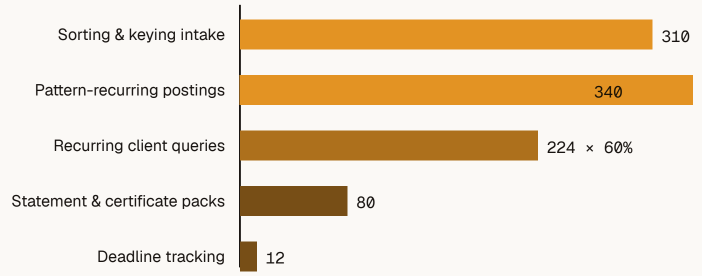

TEAM-LOGGED WEEK IN JUNE, SCALED MONTHLY AND CROSS-CHECKED IN OBSERVATION. QUERY BAR SHOWS THE RECURRING 60% SHARE OF 224 H.

> READING THE CHART
>
> The two long bars describe the same problem: documents arrive unstructured, are sorted by hand, and are then keyed by hand even when the posting matches last month's. Automating the pattern work leaves the judgment work (closings, VAT edge cases, advisory) untouched and better staffed.

28,500/mo

POSTINGS ACROSS ALL MANDATES, 71% PATTERN-RECURRING

840/mo

INBOUND CLIENT E-MAILS NEEDING AN ANSWER, Ø 16 MIN

1.8 days

Ø RESPONSE TIME; THE TOP CLIENT COMPLAINT IN YOUR 2025 SURVEY

14

DOCUMENTED DEADLINE NEAR-MISSES IN 2025; NO FINES, BUT REPEATED OVERTIME

Pain points, ranked by monthly cost

PAIN POINT	EVIDENCE	EST. COST / MONTH
Pattern postings keyed by hand	71% of 28,500 postings recur; each keyed and coded manually	≈ 340 h
Manual intake sorting & keying	3,650 docs/mo, 58% unstructured; Ø 2.1 days to booked	≈ 310 h
Recurring client queries	≈ 60% of 840 e-mails/mo answerable from the mandate file	≈ 135 h
Certificate & statement packs	Payroll certificates, interim statements assembled by hand	≈ 80 h
Deadline tracking in Excel	14 near-misses in 2025; single-point dependency on two people	≈ 12 h + risk

Hours are effort, not elapsed time. Rows do not overlap. Q1 seasonality means the annual pattern is uneven: the same work costs roughly 35% more in January–April, which is also when errors and overtime cluster.

READINESS ASSESSMENT

Readiness is 2.9 of 5, with gaps in people and governance.

Six dimensions, scored 0–5 against what the target workflow on page 7 actually requires. Scores were reviewed with both partners on July 17.

Data foundations

3.0 / 5

Abacus data is clean, and a decade of coded postings provides reference material for proposals. Unstructured client input (photos, prose) is what the intake stage is built to handle. The weak point is the three filing conventions in DeepBox.

Systems & integration

3.5 / 5

AbaConnect and Bexio APIs cover every integration the target workflow needs; DeepBox has a usable interface. The firm already operates in the cloud, so no separate infrastructure project is needed.

Process maturity

3.0 / 5

Posting and VAT routines are solid. Filing and correspondence practice varies by specialist; the Phase 1 standardization is a precondition and a benefit in its own right.

People & skills

2.0 / 5

This is the largest gap. Two senior specialists are openly skeptical ("our clients did not consent to this"); two juniors are quietly enthusiastic. The plan starts with internal drafts, nothing client-facing without a signature, and lets the results speak for themselves.

Governance & risk

2.5 / 5

Confidentiality discipline is strong culturally but undocumented for AI: no policy, and mandate agreements predate any mention of automated processing. Both are Phase 1 deliverables, completed before any client document reaches a model.

Economics

3.5 / 5

Growth is limited by capacity in a difficult hiring market; released hours convert into mandates currently being declined, or reduce unbilled Q1 overtime.

> PRECONDITIONS FOR THE BUILD PHASE
>
> **1.** One filing convention, agreed and documented (Phase 1, weeks 1–3; facilitated workshop, the team decides). **2.** AI usage policy plus a mandate-agreement annex covering automated processing, reviewed by your legal counsel (Phase 1). **3.** The two skeptical specialists named to the pilot team as reviewers with veto power. If the drafts do not convince them, that is a relevant finding in itself.

TARGET STATE

Every document is classified, then routed by confidence.

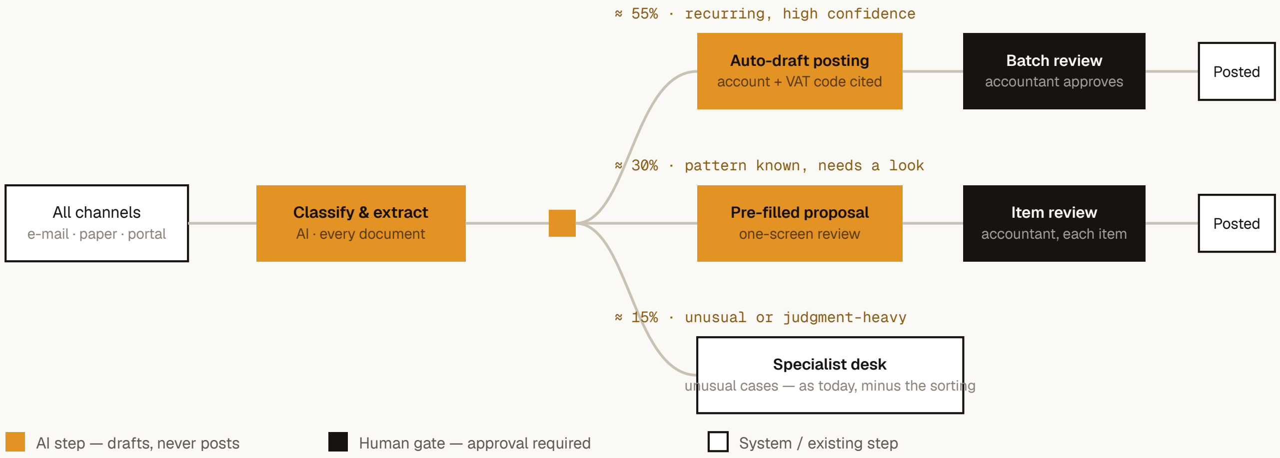

TARGET INTAKE AND POSTING LINE. SHARES ESTIMATED FROM THE 71% PATTERN-RECURRENCE MEASURED IN YOUR POSTING HISTORY; CALIBRATED IN THE PILOT.

A

Classify & extract — every document

Whatever the channel or format, each document becomes a typed record: mandate, type, period, amounts, counterparty, filed under the one agreed convention. The Sunday photo is ready for review Monday morning.

B

Posting proposals with cited evidence

Drafts carry account, VAT code, and the historical postings they are based on. High-confidence recurring items queue for batch approval; the accountant approves twenty in the time one used to take. Rejections are logged and feed back into the thresholds.

G

Human approval on every rail

Every rail ends at a person. Batch review is sampling-supported (flagged outliers first), item review is one screen, and the specialist desk is unchanged. Authorities and clients only ever see documents a person released.

C

Drafted client replies

For the recurring 60% of queries, a draft answer with sources from the mandate file appears under the e-mail. The specialist edits, signs, and sends, typically the same day instead of after 1.8 days, and the wording remains their own.

D

Deadline radar

VAT, tax, payroll, and AGM dates tracked per mandate from the systems themselves, with owner, status, and escalation. The Excel file is retired, together with the dependency on two individuals.

> DESIGN PRINCIPLE
>
> Thresholds start conservative and are adjusted per case type as override data accumulates. The 55/30/15 split is the expected end state after calibration.

SYSTEMS & THE DATA LIFECYCLE

Systems, data flow, and the Swiss-only processing boundary.

The target workflow runs on the systems you already operate — Abacus, Bexio, DeepBox, Microsoft 365 — connected rather than replaced. Every client document follows a defined lifecycle from inbox to archive. For a Treuhand, this is the page your own clients will ask about, so it is written to be shown to them.

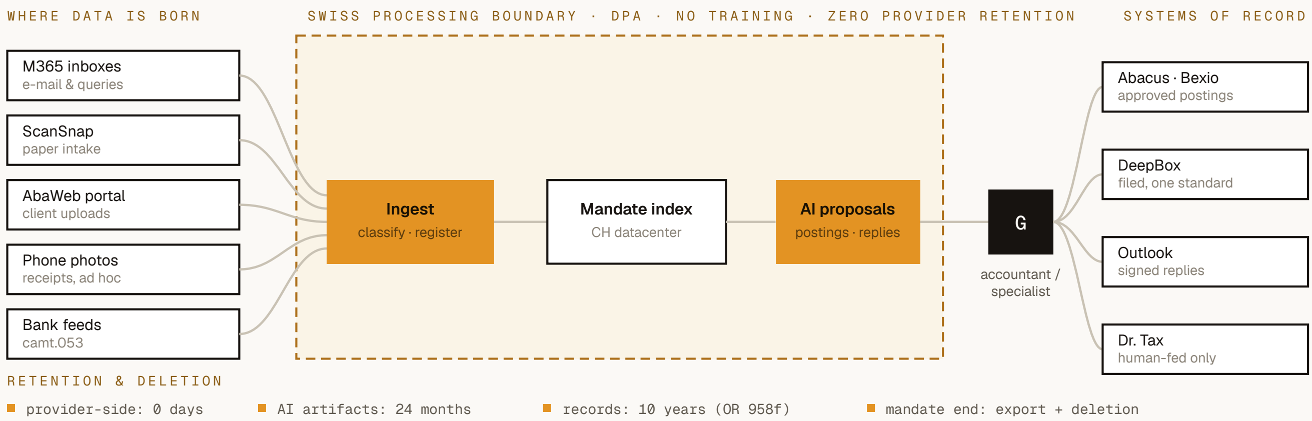

DATA LIFECYCLE IN THE TARGET WORKFLOW. ORIGINALS LIVE IN DEEPBOX AS TODAY; THE MANDATE INDEX HOLDS EXTRACTED FIELDS AND POINTERS — SCOPED STRICTLY PER MANDATE.

Your systems today and their role in the target workflow

SYSTEM	TODAY	IN THE TARGET WORKFLOW
E-mail — Microsoft 365	Shared inboxes, sorted and answered by hand	Connector reads intake and query inboxes; drafts appear under each mail; sending stays in Outlook, signed
Accounting — Abacus + AbaWeb	Postings keyed manually; history unused	AbaConnect writes approved postings; ten years of coded history feeds the proposals
Small mandates — Bexio	Separate manual flow	Same proposal engine via the Bexio API; one workflow serving two backends
Archive — DeepBox	Three naming conventions coexist	One standard; the AI files every document; originals never leave the box
Tax — Dr. Tax	Stand-alone	Outside AI scope by design; receives human-approved data only
Scanning — ScanSnap	Paper piles await sorting	Hot-folder scans auto-registered on arrival, like every other channel

Lifecycle rules

- **Capture:** every document registered per mandate on arrival, with period and source.
- **Processing:** Swiss-hosted inference only, under DPA; zero provider retention; no training on client data.
- **Access:** mirrors Abacus mandate permissions; payroll data isolated to the payroll team, as today.
- **Consent:** plain-language processing annex with the annual engagement letters; clients who decline stay on the manual path.
- **Retention & deletion:** the 10-year record duty (OR 958f) is untouched; AI artifacts expire after 24 months; ending a mandate triggers export plus scheduled deletion of its index.

EXCLUDED BY DESIGN

Training on client data · any processing outside Switzerland · client documents in free public AI tools · cross-mandate visibility: every AI request is scoped to exactly one mandate, enforced technically as well as by policy.

WHAT CLIENTS WILL NOTICE

Documents are booked the next morning and replies arrive the same day, with a one-page annex that explains the processing in plain language.

OPPORTUNITY PORTFOLIO

Of six opportunities identified, five are recommended now.

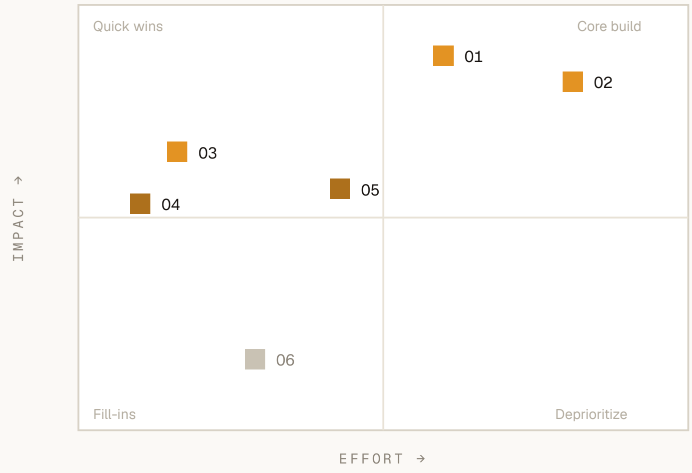

IMPACT = CAPACITY RELEASED + CLIENT-VISIBLE SPEED · EFFORT = BUILD + CHANGE COST.
POSITIONS REVIEWED WITH THE PARTNERS.

PORTFOLIO LOGIC

O3 and O4 go first: they bring visible relief within weeks, and nothing client-facing leaves without a signature. The skeptics sit in the review seat from day one. O1 and O2 are one system; the intake store feeds the proposal engine.

Deprioritized

O6 — meeting-note to advisory memo drafting. Attractive, but it touches the most sensitive conversations and the least structured input. Revisit once the team trusts the system on lower-stakes work.

#	OPPORTUNITY	WHAT IT REPLACES	CAPACITY / MO	EFFORT	PHASE
01	**Intake classification & extraction**	Manual sorting and keying of 3,650 docs/mo across five channels	≈ 190 h	Medium	2
02	**Posting proposal engine**	Hand-keying the recurring 71% of ≈ 28,500 postings	≈ 170 h	High	2
03	**Drafted client replies**	Composing recurring answers from scratch, interrupt-driven	≈ 80 h	Low	1
04	**Deadline radar**	The Excel file and two people's memories	≈ 10 h + risk	Low	1
05	**Certificate & statement packs**	Hand-assembled payroll certificates and interim statements	≈ 45 h	Medium	3
06	Advisory memo drafting	Writing up client meetings after hours	n/a	Medium	Later

Capacity figures assume full adoption at calibrated thresholds. Σ O1–O5 ≈ 495 h/mo gross; the plan assumes ≈ 280 h/mo to allow for review time at the gates, threshold conservatism in year one, and Q1 lumpiness.

ROADMAP & INVESTMENT

The plan runs 22 weeks, built around the closing season.

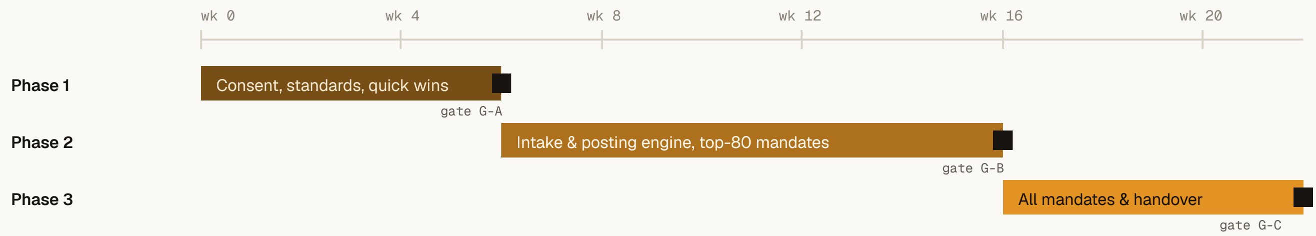

Phase 1 — Consent, standards, quick wins

WEEKS 1–6

One filing convention agreed; AI policy and mandate-agreement annex through legal review; O3 (drafted replies) and O4 (deadline radar) live for the pilot group including both skeptics.

Gate G-A: policy signed; ≥ 70% of drafted replies sent with light edits; deadline radar has replaced the Excel.

Phase 2 — Intake & posting engine

WEEKS 6–20

O1 and O2 built on the standardized store; confidence routing calibrated on the 80 largest bookkeeping mandates against six months of history before touching live work.

Gate G-B: ≥ 90% classification accuracy on backtest; ≥ 45% of live postings auto-drafted and approved in batch; accountant override rate < 20% and falling.

Phase 3 — All mandates & handover

WEEKS 20–22 + Q1 STANDBY

Rollout to the remaining mandates and O5 certificate packs; runbook, admin training, and threshold governance handed to the internal owner. Everything that changes daily routines is live and stable before January; what remains in the new year is extending the proven setup to the smaller mandates, with standby support on call through the first closing season.

Gate G-C: Alpstein operates two weeks without external help; Q1 readiness checklist signed by both partners.

Indicative budget

ITEM	CHF
Phase 1 incl. policy work, filing standard, both quick wins	32,000
Phase 2 range narrows once Gate G-A results are in	60,000–75,000
Phase 3 incl. Q1 standby support	30,000–40,000
Running costs / month CH-hosted models, licenses, monitoring	1,100–1,600

Estimates at typical Swiss market rates for a build of this scope, whoever executes it — an external integrator or internal IT with targeted support. The ranges are driven mainly by integration depth (Abacus, Bexio, DeepBox) and by how much threshold calibration the pilot reveals.

> PAYBACK, CONSERVATIVE CASE
>
> Plan against ≈ 280 h/mo at 2/3 realization in year one: ≈ 187 h/mo × CHF 105 ≈ **CHF 235,000/yr**, against a total build of CHF 122,000–147,000 → **payback in 7–9 months**, before any new mandates taken on with the freed capacity are counted.

KPIs tracked from week 1

- Arrival-to-booked lag (days); share auto-drafted vs. keyed
- Classification accuracy; accountant override rate per rail
- Ø client response time; share of replies sent from drafts
- Q1 overtime hours vs. 2025/2026 baseline

GOVERNANCE, RISK & DATA PROTECTION

Confidentiality requirements and how the design meets them.

Data protection & confidentiality

- **Hosting:** model inference exclusively on Swiss-hosted infrastructure under a data processing agreement; no client data used for provider training. This is stricter than the revised FADP requires, and appropriate for a fiduciary mandate.
- **Client consent:** mandate agreements gain a plain-language annex on automated processing support, rolled out with the annual engagement letters. Legal counsel reviews the annex before go-live. Clients who decline stay on the manual path; the workflow supports both.
- **Access:** per-mandate permissions mirror Abacus exactly. Payroll data, the most sensitive category, stays within the payroll team's scope, as today.
- **Retention & audit:** every classification, proposal, and human decision is logged per mandate. Who saw which document, and why, becomes easier to answer than it is today.

Human oversight policy

- No document is posted, filed, or sent without a named person's approval. Batch review remains a genuine check: sampling puts flagged outliers first.
- Client-facing text is always signed by a specialist. The client receives a personally signed reply, and accountability stays with the signer.
- Submissions to authorities (VAT, tax, social insurance) remain fully human at the final step, unchanged.

Key risks and how the design mitigates them

RISK	MITIGATION
A wrong posting slips through batch approval	Conservative thresholds at start; outlier-first sampling; monthly reconciliation unchanged; every auto-draft reversible with one click and logged.
A drafted reply contains a wrong answer	Drafts cite their sources from the mandate file; specialists sign every send. Statements without a source are flagged in the draft.
Client trust damage from disclosure	Proactive transparency: the annex explains what is automated and what is not. In our experience, clients receive the faster turnaround as a clear benefit once it is explained.
Team rejection	Skeptics hold review seats with veto from week 1; high override rates are investigated as findings; no productivity targets are tied to the system during the pilot.
Vendor lock-in	Alpstein owns the standardized store, prompts, and thresholds; export is part of the standard handover. Ongoing support can sit with any provider, or in-house.

REGULATORY NOTE

The workflow makes no automated decisions with legal effect for individuals (FADP Art. 21 stays untriggered): every posting, filing, and reply is human-approved. Payroll processing continues under existing social-insurance obligations; the AI only assembles drafts from data you already hold.

HOW TO PROCEED

Timing, roles, and what to require of whoever builds.

Recommended sequence

WHEN	WHAT
Week of Aug 10	Walkthrough of this report with both partners, then with the full team (same content for both groups). Decision on Phase 1.
Week of Aug 24	Implementation route chosen: external integrator, or internal IT with targeted support. Legal counsel briefed on the annex; filing-standard workshop scheduled; pilot group named, skeptics included.
Sep 7	Phase 1 kickoff in St. Gallen. Deadline radar live by end of September; drafted replies in the pilot inboxes by mid-October.
Mid-Oct	Gate G-A review: continue into the core build, adjust, or stop with a working radar, a signed policy, and one filing standard either way.

Roles the plan requires

- An internal product owner at ~20% (proposed: the senior accountant)
- AbaConnect and DeepBox access; posting history export
- Pilot group availability: 30-min weekly stand-up
- Legal counsel time for the annex review

No-regret first steps

- Record the arrival date on every intake document from August; the Ø 2.1-day lag becomes measurable
- Start the filing-standard discussion now; one DeepBox convention pays off with or without AI
- Put the consent annex on legal counsel's desk early; it gates everything else in the plan

What to require from any implementation partner

- Swiss-only hosting for model inference under a DPA, as specified on page 8; zero provider retention
- Integration references covering Abacus/AbaConnect, Bexio, and DeepBox, not generic AI projects
- The acceptance gates G-A to G-C written into the agreement, with a stop right at each gate
- All measurements published in full, override rates included; training for the pilot group and admins
- Full handover: runbook, admin rights, exportable data, no lock-in

VALIDITY OF THIS ASSESSMENT

Figures and estimates are valid until October 31, 2026; after that date, re-check the volume baseline before committing budget. The timeline is built backward from one hard constraint: nothing new goes live during closing season. A decision after mid-September pushes the core build past January and costs a year of Q1 relief.

THE KEY FIGURE

Of every ten postings the team keys, **seven follow a pattern the system can draft**. The accountant's work shifts from typing entries to reviewing and approving them, which matters most in the January–April closing season.

Leitspur
SME AI SYSTEMS

Practical AI, engineered into daily work.

Leitspur builds AI systems for small and mid-sized companies — workflow by workflow, with human review built in and measurable results as the only definition of done.

CONTACT
info@leitspur.ai

WEB
leitspur.ai

REFERENCE
LS-SAMPLE-02 · v1.0