




■ AI READINESS & WORKFLOW MAPPING · SAMPLE REPORT

AI readiness and a **target workflow** for the quote-to-order process.

PREPARED FOR

Berger Präzisionsmechanik AG

Bellach, Solothurn · July 2026

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

REFERENCE

LS-SAMPLE-01

FIELDWORK

June 30 – July 11, 2026

ISSUED

July 24, 2026 · v1.0

STATUS

Sample report — free to share

ABOUT THIS ENGAGEMENT

Scope, method, and structure of this report.

Berger Präzisionsmechanik AG asked Leitspur to map its quote-to-order process end to end, assess where the company stands on AI readiness, and design a target workflow in which the repetitive steps are automated and engineers retain every approval decision. The plan is written to stand alone: internal IT, any competent integrator, or Leitspur can execute it, and its requirements and gates apply whoever builds.

Method

- 11 structured interviews: sales administration, AVOR / work preparation, quality, production planning, and management.
- Two days of on-site observation in Bellach, following live RFQs from inbox to signed quote.
- Systems walkthrough: Abacus ERP, the NAS drawing archive, Excel cost sheets, and the quotation letter templates.
- Quantitative baseline from the RFQ log, January–June 2026 (n = 447 requests), cross-checked against 60 sampled quote files.

Scope

In scope: RFQ intake through order confirmation; first-article inspection (FAI) documentation. **Out of scope by agreement:** shop-floor scheduling, machine data, and predictive maintenance; noted where relevant but not assessed.

HOW FIGURES WERE DERIVED

Volumes come from your systems; touch times were measured during shadowing and validated in interviews. Where we estimate, we say so and round conservatively. All franc figures use a loaded internal rate of CHF 92/h agreed with management.

Contents

01	Executive summary	03
02	The quote-to-order process today	04
03	Where the hours go	05
04	AI readiness assessment	06
05	The target workflow	07
06	Systems & the data lifecycle	08
07	Opportunity portfolio	09
08	Roadmap and investment	10
09	Governance, risk, data protection	11
10	How to proceed	12

Company snapshot

Company	Berger Präzisionsmechanik AG, Bellach SO — CNC turning and milling, founded 1968, family-owned
People	72 employees, of which 2 senior AVOR planners carry most quoting
Markets	Medtech, machine building, automotive suppliers (CH / DE / FR)
Systems	Abacus ERP (on-prem), NAS drawing archive (~18,000 part records), Excel costing, Word letters
Certification	ISO 9001, ISO 13485 (medtech line)

EXECUTIVE SUMMARY

Quoting takes **4.5 days** on average; customers typically decide within two.

74 RFQ/moAVERAGE REQUEST VOLUME,
JAN-JUN 2026**4.5** daysMEDIAN QUOTE TURNAROUND
(WORKING DAYS)**3.4** hTOUCH TIME PER QUOTE,
MEASURED**26** %

WIN RATE ON QUOTED RFQS

What we found

- Quoting consumes **≈ 252 hours per month**, most of it from your two senior AVOR planners, the same people who prepare production. Quoting and work preparation compete for the same two people.
- More than half of quote touch time is **searching and re-typing**: finding similar parts on the NAS and in memory (1.1 h) and structuring RFQ data from PDFs and e-mail (0.6 h) before real engineering judgment begins.
- In peak months, **14% of RFQs lapse unanswered**. Two key customers now award standard parts to the first credible quote, typically within 48 hours.
- The knowledge that makes quotes accurate (which part is comparable, which machine, which cycle time) is **held by two people** and an unsearchable folder tree.
- FAI documentation for the medtech line repeats the same pattern: 22 reports a month, 2.1 hours each, largely re-keying CMM output into Word.

What we recommend

- Build an **AI-assisted quoting line**: automatic RFQ extraction, similar-part retrieval over your 18,000 historical records, and a drafted cost sheet, with the engineer approving every quote before it leaves.
- Start with two **low-risk quick wins** (quote letter generation, FAI report drafting) to build trust while the retrieval foundation is prepared.
- Sequence in three phases over ~20 weeks; each phase has a measurable exit gate and can stand alone if you stop after it.

READINESS VERDICT

2.9 / 5, conditionally ready. Process discipline and economics are strong; data consolidation and an AI usage policy are the two gaps to close before building. Neither is a blocker; both are Phase 1 work.

Expected effect at full rollout

METRIC	TODAY	TARGET
Quote turnaround, standard RFQs	4.5 d	< 1 d
Touch time per quote	3.4 h	≈ 1.5 h
RFQs lapsing unanswered	14 %	≈ 0 %
Capacity released, all workflows	—	≈ 175 h/mo

At CHF 92/h loaded, ≈ 175 h/mo corresponds to a gross capacity value of roughly CHF 190,000 a year. We plan on the conservative case, about two thirds of that realized in year one, and still see payback inside twelve months (p. 10).


HEARD ON SITE

"If the drawing lands on the right desk, we're fast. If it has to wait for Markus or me, it waits." — AVOR planner

CURRENT STATE

Current quoting process: eight steps across three desks.

The sequence below is the standard path of an RFQ through Berger today, drawn from 60 sampled quote files and two days of live observation. Most of the elapsed time is spent in two queues, both in front of the same two planners.



WHAT THE MAP SHOWS

Of 4.5 elapsed days, only ≈ 3.4 hours are actual work. Two waits — the AVOR queue (1.4 d) and management sign-off (0.6 d) — plus batching of steps 1–2 account for over 80% of turnaround. Follow-up calls happen for fewer than 1 in 5 quotes.


Observations relevant to the design

- **Intake is untyped.** RFQs arrive as e-mail prose, PDF drawings, STEP files, occasionally photographed sketches. Nothing is machine-readable until a person re-types it.
- **The similar-part search is the core skill.** Planners navigate the NAS by customer and year, then rely on memory. New hires cannot quote independently for 12–18 months mainly because of this step.
- **Costing is consistent.** The Excel cost sheet is disciplined and audit-friendly, a good foundation. The AI should pre-fill it rather than replace it.
- **Nothing feeds back.** Won/lost outcomes stay in Abacus; the cost sheets never learn. Win-rate by part family is unknown today.

QUANTIFIED BASELINE

More than half of the touch time goes to search and re-typing.

Touch time per quote by activity — measured, hours



Σ 3.4 H · SHADOWING SAMPLE N = 14 QUOTES, CROSS-CHECKED IN INTERVIEWS

READING THE CHART

The two bright bars, searching for comparable parts and building the cost sheet, are where AI assistance pays most. The search is automatable retrieval; the cost sheet can be pre-filled from the retrieved parts. Engineering judgment remains in the process: the target is a reviewed 1.5 h per quote.

252 h/mo

TOTAL QUOTING EFFORT (74 RFQ × 3.4 H)

CHF 278k/yr

COST OF THE CURRENT PROCESS AT CHF 92/H

46 h/mo

FAI REPORTS: 22 × 2.1 H, MEDTECH LINE

20 h/mo

ORDER CONFIRMATIONS KEYED INTO ABACUS (48 × 25 MIN)

Pain points, ranked by monthly cost

PAIN POINT	EVIDENCE	EST. COST / MONTH
Similar-part search from memory	1.1 h per quote; blocks new hires from quoting independently	≈ 81 h
Manual cost-sheet assembly	1.2 h per quote; values re-derived that exist in past sheets	≈ 89 h
Untyped RFQ intake	0.6 h per quote re-typing PDFs/e-mail into Excel and folders	≈ 44 h
FAI report writing	Re-keying CMM measurement output into Word protocols	≈ 46 h
Lapsed RFQs in peak months	14% unanswered (Mar & May 2026); revenue foregone, not costed here	—

Hours are effort, not elapsed time, and overlap is removed (the 252 h quoting total already contains the first three rows). We deliberately do not monetize lapsed RFQs, because win-rate data by part family is not yet reliable enough. We count this as upside rather than as part of the business case.

READINESS ASSESSMENT

Readiness assessment: **2.9 / 5**, conditionally ready.

We score six dimensions on a 0–5 scale, anchored to what the target workflow on page 7 actually requires, rather than to an abstract maturity model. Scores were reviewed with management on July 11.

Data foundations

2.5 / 5

18,000 part records exist but sit in a folder tree keyed by customer and year. Quote history is complete in Excel yet unlinked to outcomes. Consolidation is a well-bounded Phase 1 task.

Systems & integration

3.0 / 5

Abacus exposes a usable API (AbaConnect); the NAS is well organized and access-controlled. There is no cloud estate yet; the hosting decision is open and is made in Phase 1.

Process maturity

3.5 / 5

ISO-certified discipline shows: quoting steps are consistent across both planners, and the cost sheet is a de-facto standard. Automation lands well on stable processes like this one.

People & skills

2.5 / 5

AVOR is openly in favor: the planners want the search and typing off their desks. AI experience is limited to private ChatGPT use; structured onboarding and floor-level champions are needed.

Governance & risk

2.0 / 5

No AI usage policy; customer NDAs (notably both medtech OEMs) restrict where drawings may be processed. Resolvable with CH/EU-hosted models and a data-handling annex, but it must be settled before the first drawing is processed externally.

Economics

4.0 / 5


High hourly value, clear cost of delay, and a bottleneck on revenue-critical work. Few SME processes justify investment as directly as this quoting line.

PRECONDITIONS FOR THE BUILD

1. Quote history and part records consolidated into one indexed store (Phase 1, weeks 1–4). 2. AI usage policy signed, including the customer-NDA review for the five largest accounts (Phase 1, parallel). 3. A named product owner on Berger's side with ~20% availability. We propose the junior AVOR planner: reviewing at the gate doubles as quoting training.

TARGET STATE

The target quoting workflow and its two approval gates.



TARGET QUOTE LINE. COST SHEET, TEMPLATES, AND SIGN-OFF RULE ARE UNCHANGED; SEARCH, STRUCTURING, AND DRAFTING MOVE TO THE ASSISTANT.

2 Extract & structure the RFQ

Reads e-mail, PDF drawings, and STEP metadata; produces a typed request record (material, dimensions, tolerances, quantities, deadline) and files everything correctly. Flags missing information and drafts the clarification e-mail instead of waiting a day for a planner to spot the gap.

3 Similar-part retrieval

Searches all 18,000 historical part records by geometry class, material, tolerance band, and customer, and returns the five closest parts with their actual cost sheets and outcomes. Part knowledge that today lives in the planners' memory becomes a lookup available to everyone.

4 Cost sheet draft

Pre-fills your existing Excel structure from the retrieved parts: machine group, setup and cycle estimates, material at current rates, margins per your rules. Every pre-filled cell cites its source part, so the planner verifies values instead of re-deriving them.

G1 Engineer review gate

A planner reviews and adjusts the draft with full authority, on one screen, with all sources visible. Nothing proceeds without this approval. The gate is also where the system learns from corrections; override rates are measured during the pilot.

5 Quote letter draft

Generates the letter in Berger's format and language conventions, both German and French customers covered, prices from the approved sheet only.

G2 Release & sign-off

Existing authority rule unchanged: quotes above CHF 20,000 are countersigned by the sales lead. Below the threshold, the reviewing engineer releases directly.


DESIGN PRINCIPLE

The gates are not a transition phase to be automated away later. A quote is a price commitment and remains a human decision. What changes is what that person spends their time on.

SYSTEMS & THE DATA LIFECYCLE

How data moves through your existing systems.

Nothing in the target workflow replaces a system you run today. What changes is that every document follows a defined lifecycle: where it originates, where it is processed and stored, who approves it, and when it is deleted. This page sets out where your data goes at each step.



DATA LIFECYCLE IN THE TARGET WORKFLOW. SYSTEMS OF RECORD ARE UNCHANGED; THE INDEX STORES EXTRACTED FIELDS AND POINTERS RATHER THAN COPIES OF THE ARCHIVE.

Current systems and their role in the target workflow

SYSTEM	TODAY	IN THE TARGET WORKFLOW
E-mail — Microsoft 365	RFQs arrive in a shared mailbox, triaged by hand	Ingest connector reads the RFQ mailbox; letters still go out from Outlook
Files — Synology NAS	18,000 part folders, permissions by team	Stays the system of record for drawings; the index references files without moving them
ERP — Abacus, on-prem	Quotes and orders keyed by hand	AbaConnect API: quote entry written after release; win/loss read back nightly
Costing — Excel	The de-facto cost-sheet standard	Format unchanged; pre-filled by the assistant, every cell citing its source
QC — CMM software	Measurement exports re-typed into Word	Exports parsed directly into FAI report drafts
Backup — IT partner	Encrypted offsite, NAS + ERP	Indexed store joins the existing backup plan; no separate regime

Lifecycle rules

- **Capture:** every RFQ registered on arrival with customer, source, and timestamp; the log is generated as a by-product.
- **Processing:** inference on CH/EU infrastructure under DPA; zero provider retention; no training on your data.
- **Storage:** extracted fields and pointers in the indexed store (CH region or on-prem beside the NAS; Phase 1 decision); drawings stay where they are.
- **Access:** the index inherits NAS and Abacus permissions unchanged; nothing becomes more visible than it is today.
- **Deletion:** remove a file from the NAS and the index forgets it within 24 hours.

EXPLICITLY EXCLUDED


Training provider models on your data · processing outside CH/EU · customer drawings in free public AI tools · bulk copies of the archive. The index holds fields and pointers; your NAS remains the only full copy.

COMMERCIAL RELEVANCE

Under both medtech NDAs, data handling is also a sales argument: after Phase 1, the question of where a customer drawing is processed has a precise, documented answer. Few suppliers can currently offer that.

OPPORTUNITY PORTFOLIO

Six opportunities, of which five are recommended now.



PORTFOLIO LOGIC

O4 and O5 come first deliberately: both are small and low-risk, and they let the team get used to reviewing AI drafts. O1–O3 are three stages of one system and share the indexed part store built in Phase 1.

Deprioritized

O6 — shop traveler & NC note search. Real demand from the floor, but it depends on digitizing 20 years of handwritten travelers. Revisit after Phase 3 with scanning cost data.


IMPACT = CAPACITY RELEASED + TURNAROUND EFFECT · EFFORT = BUILD + CHANGE COST.
POSITIONS REVIEWED WITH MANAGEMENT.

#	OPPORTUNITY	WHAT IT REPLACES	CAPACITY / MO	EFFORT	PHASE
01	RFQ intake extraction	Re-typing PDFs and e-mail into Excel; drafts clarification e-mails	≈ 40 h	Medium	1–2
02	Similar-part retrieval	NAS navigation and personal memory; de-risks succession	≈ 70 h	High	2
03	Cost-sheet draft assistant	Manual re-derivation of setup, cycle, and material lines	≈ 45 h	Medium	2
04	Quote letter + ERP entry	Word formatting and manual Abacus entry after approval	≈ 20 h	Low	1
05	FAI report drafting	Re-keying CMM output into Word protocols (medtech line)	≈ 26 h	Low	1
06	Traveler / NC note search	Asking the veteran machinist	n/a	High	Later

Capacity figures assume full adoption and are not additive with turnaround gains (same hours, seen two ways). Σ O1–O5 ≈ 201 h/mo gross; we plan against ≈ 175 h/mo to allow for review overhead at the gates.

ROADMAP & INVESTMENT

Implementation plan: three phases over 20 weeks.



Phase 1 — Foundations & quick wins

WEEKS 1-6

Consolidate quote history + part records into one indexed store. Ship O4 (letter + ERP entry) and O5 (FAI drafts) to three users. AI policy and customer-NDA review completed.

Gate G-A: extraction ≥ 95% field accuracy on a 100-RFQ backtest; both quick wins in daily use; policy signed.

Phase 2 — Core quoting assistant

WEEKS 5-14

Build O1-O3 on the Phase 1 store: retrieval, cost-sheet drafting, review gate UI, Abacus integration. Pilot with both planners on live RFQs in parallel to the old path.

Gate G-B: ≥ 60% of standard RFQs quoted in < 1 day over four consecutive weeks; planner override rate < 25% and falling.

Phase 3 — Scale-out & handover

WEEKS 14-20

All quoting users onboarded; FAI drafting extended to all report types; runbook, admin training, and governance handbook handed to the internal owner. External support drops to on-call.

Gate G-C: KPI dashboard live; Berger operates the system two weeks without external help.

Indicative budget	
ITEM	CHF
Phase 1	28,000
incl. data consolidation & both quick wins	
Phase 2	55,000–70,000
scope is known at Gate G-A; re-estimate there	
Phase 3	25,000–35,000
Running costs / month	900–1,400
CH/EU-hosted models, licenses, monitoring	

Estimates at typical Swiss market rates for a build of this scope, whether it is delivered by an external integrator or by internal IT with targeted support. The ranges are driven mainly by the depth of the Abacus integration and the review-gate interface.

PAYBACK IN THE CONSERVATIVE CASE

Assume only two thirds of the ≈ 175 h/mo materializes in year one: ≈ 117 h/mo × CHF 92 ≈ **CHF 129,000/yr** against a total build of CHF 108,000–133,000 → **payback in 10–12 months**, before any effect on win rate or lapsed RFQs.

- KPIs tracked from week 1**
- Median quote turnaround (days) and touch time (h)
 - Extraction field accuracy; retrieval hit rate (top-5)
 - Planner override rate at Gate G1 (trust indicator)
 - Share of RFQs lapsing unanswered

GOVERNANCE, RISK & DATA PROTECTION

How customer drawings and quoting data are protected.

Data protection & confidentiality

- **Hosting:** all model inference on CH- or EU-region infrastructure under a data processing agreement; no customer data used for provider model training. This satisfies the revised FADP and, for EU customers, GDPR expectations.
- **Customer NDAs:** the five largest accounts' NDAs are reviewed in Phase 1. Both medtech OEM agreements permit subcontracted data processing under DPA; written confirmation is a Phase 1 deliverable, obtained before any drawing is processed.
- **Access:** the part-record index inherits NAS permissions; quoting data stays inside the quoting team, as today.
- **Retention & audit:** every AI draft, retrieval result, and human decision is logged. If a customer asks how a price was derived, the derivation can be retrieved from the log.

Human oversight policy

- No quote, e-mail, or ERP entry leaves the company without a named person's approval. AI output is a draft by definition.
- Price authority is unchanged (sign-off > CHF 20,000). The AI never adjusts margins; it applies your rules and cites its sources.
- The manual path stays available and documented. If the system is down, quoting continues manually as it does today.

Key risks and how we mitigate them

RISK	MITIGATION
Extraction errors reach a quote	Gate G1 review of every field with source shown; accuracy tracked per field; fields below 98% confidence are visually flagged for checking first.
Planners silently ignore the drafts	Override rate is a first-class KPI, reviewed in the weekly pilot stand-up. A persistently high override rate triggers fixes to the system rather than pressure on the planners.
Retrieval suggests a bad comparable	The system returns the five closest parts with similarity rationale, and the planner selects. Wrong suggestions are dismissed in one click and logged for tuning.
Key-person risk shifts to the vendor	Phase 3 hands over runbook, admin rights, and exportable data. Berger owns the index and prompts; ongoing support is optional.
Model or provider change	Architecture is model-agnostic behind one interface; switching providers is a configuration exercise, re-validated against the backtest set.

REGULATORY NOTE

The workflow involves no automated decisions about individuals; FADP exposure is limited to business contact data in RFQs. ISO 13485 documentation duties are unaffected: FAI reports remain reviewed and signed by quality staff; the AI only assembles the draft from CMM output.

HOW TO PROCEED

Choosing the route by mid-August allows a September 1 start.

Proposed sequence

WHEN	WHAT
Week of Aug 4	Walkthrough of this report with management and both AVOR planners (90 min, on site). Decision on Phase 1 and the build route.
Week of Aug 18	Implementation partner selected or internal owners named; Abacus API user and NAS read access provisioned; NDA review letters sent to the five key accounts.
Sep 1	Phase 1 kickoff in Bellach. First quick win (quote letters) live for pilot users by end of September.
Mid-Oct	Gate G-A review: continue, adjust, or stop, with a working quick win and a consolidated data store either way.

Roles the plan requires

- Product owner at ~20% (proposed: junior AVOR planner)
- Abacus API access, NAS read access, sample exports
- Both planners for weekly 30-min pilot stand-ups
- Decision authority at each gate

No-regret first steps

- Timestamp every RFQ on arrival and at quote release; the log becomes the baseline the gates are measured against
- Record won/lost outcomes in Abacus consistently, so win rate by part family becomes measurable
- Keep customer drawings out of free public AI tools until the usage policy is signed

What to require from any implementation partner

- All processing on CH/EU infrastructure under DPA, per the boundary defined on page 8
- Integration references for Abacus, Microsoft 365, and NAS-based archives
- Acceptance gates with stop rights; measurements published in full
- Training for all quoting users and the internal admin
- Full handover: runbook, admin rights, exportable data, no lock-in

VALIDITY OF THIS ASSESSMENT

Figures and estimates in this report are valid until October 31, 2026. The baseline (447 RFQs, H1 2026) ages well, but if volumes shift materially (for example if the rumored medtech frame contract lands), Phase 2 sizing should be re-checked at Gate G-A before committing further.

THE DAILY EFFORT

Every working day, roughly **12 hours** of skilled AVOR time goes into searching, re-typing, and formatting. The target workflow returns most of that time to engineering and lets customers receive answers the same day.



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-01 · v1.0