



OVERVIEW

Before you start: Readiness Check

Use this guide for one AI use case at a time. Before launching a pilot, use the checklist below to see whether your company is ready.

Key point: AI is easy to start, but difficult to control and scale. The three levels below address that.

LEVEL 1

Strategic



Start from a real problem, not the technology. Pick one small, measurable use case.

→ Page 2

LEVEL 2

Governance



Set data limits, approved tools and human oversight before scaling.

→ Page 3

LEVEL 3

Provider



Choose tools through risk-based trade-offs, not hype or popularity.

→ Page 4

Quick readiness check

Tick what is already true. The unticked items show the missing basics to fix before moving forward.

These nine items map to the most common failure points found across the practitioner interviews behind this guide.

☐

One person is accountable for this AI use case.

☐

We know which AI tools employees currently use.

☐

We have written rules on which data this use case may send to AI tools.

☐

Our planned use case targets a concrete, repetitive task.

☐

We can measure the gain (time, cost, errors).

☐

We know where our data would be processed (CH / EU / US).

☐

A human reviews AI outputs before they reach customers.

☐

The data this use case needs is structured and reasonably clean.

☐

Management agrees to support and review this pilot.

0–3 checked

Do not launch a pilot yet. Fix the missing basics first.

4–6 checked

Pilot only a low-risk use case with limited data exposure.

7–9 checked

Ready to define your use case on Page 2.

Stop here if:

Nobody owns or supports this use case, you don't know which AI tools employees currently use, or there is no clear rule for sensitive data.

Safety rule:

If data rules, tool visibility or human review are missing, don't start with sensitive or customer-facing use cases.

Need external input?

Before committing to a vendor, start with neutral sources: your industry association, universities or research institutes such as HSLU, ETH or EPFL, vendor-neutral AI communities such as the Digital & AI Community in Lucerne, or independent advisors with references in similar use cases.



LEVEL 1 - STRATEGIC

Scope your Use Case

Fill out this page before choosing any AI tool. If you cannot define the problem, baseline, testable version and success metric, stop: the use case is not ready.

01. USE CASE NAME
e.g. "Auto-extract data from supplier invoices"
.....

02. OWNER
One named person
.....

03. 30-DAY DEADLINE
DD.MM.YYYY
.....

04. THE PROBLEM
A repetitive, slow or error-prone task. State the pain point, not the technology.
One sentence is required.
.....

05. BASELINE
Current time, cost, volume or error rate.
A number is required.
.....

06. SMALLEST TESTABLE VERSION
What is the smallest version you can test on real work within 30 days? Define what AI does and what the human still verifies.
.....

07. SUCCESS METRIC
A measurable target, e.g. "reduce time per invoice by 40% while keeping errors below 3%."
.....

08. DATA SENSITIVITY
Public / Internal / Confidential / Highly sensitive.
Confirm on Page 3.
.....

09. STOP RULE
If the pilot is not on track to meet the success metric by day ____, we stop.
.....

10. MAIN RISK TO TEST
What is most likely to make this pilot fail: data quality, tool accuracy, user adoption, integration effort, cost or something else?
.....

"The question should not be: what can we do with AI? It should be: what can we solve with AI?"

Ronny Fischer - AI Initiative Lead, InfoGuard



LEVEL 2 - GOVERNANCE


Where can your Data go?

Before choosing a tool, classify the data this use case will use.
Then decide what level of control it requires.

1. Classify the Data

TIER	EXAMPLES	DEFAULT AI RULES
PUBLIC	Public reports, public website text, public FAQs, general market information	Any approved AI tool. Still review outputs before publication.
INTERNAL	Internal notes, draft procedures, non-sensitive meeting summaries	Approved company tools only. No private accounts.
CONFIDENTIAL	Contracts, financial documents, customer data, source code, business strategy	Approved CH/EU-hosted or contractually cleared tool. Clarify data-processing terms, access control, training use and retention before testing.
HIGHLY SENSITIVE	Health data, therapy notes, biometric/genetic data, credentials, trade secrets, security-sensitive information, legally regulated data	Do not use external AI by default. Use only controlled/self-hosted or expert-approved setups.

2. Decide how to Proceed



Keep in mind:

Risk does not only depend on what enters the tool, but also on how the result is used, especially for customers, decisions or official records.



There is no universal "best" provider. Use your data tier and use-case needs to compare the trade-offs below.

Provider Type	Data Location / Jurisdiction	Capability & Workflow Fit	Contract & Control	Cost Model	Lock-in Risk*
Global AI platforms e.g. OpenAI, Anthropic, Google	US / Global	Top-tier, easy	Standard terms	Usage-based	High
CH / EU-hosted or sovereign cloud e.g. Infomaniak, AlpineAI, Swisscom	CH / EU	Strong models, more setup	Negotiable	Higher baseline	Medium
Self-hosted e.g. open models on company-controlled infrastructure	On-premise	Tailored, technical	Direct control	High upfront	Low
SaaS tools with built-in AI e.g. Copilot, Gemini, Zendesk AI, HubSpot Breeze	Vendor-set	Task-specific, embedded	Bundled	Predictable	High

Control has a cost:

IF YOUR PRIORITY IS...

Sensitive or regulated data	➔	CH/EU-hosted, sovereign cloud, or expert-approved setup
Fast deployment	➔	SaaS tools with built-in AI or global AI platform
Workflow adoption	➔	A tool employees already use
Avoiding lock-in	➔	Easy data export, portable workflows, realistic provider switching
Predictable cost	➔	Fixed-price SaaS or clear usage limits

- 1 Limited portability of the workflow
- 2 No usable export of data, prompts or outputs
- 3 Dependence on one specific model version
- 4 Unpredictable costs at higher usage
- 5 Unclear terms on retention, training use or subcontractors

Write down the 1–2 providers you will evaluate further. If you cannot fill at least one row, you don't yet have a shortlist. Go back to the provider matrix above.

PROVIDER	WHY IT FITS THIS USE CASE	REQUIRED CHECKS BEFORE PILOT
Provider 1...	Match against your data tier, workflow needs, cost limits...	Verify terms, data location, training use, retention, export, costs...
Provider 2... (optional)		

Marc Pouly - AI Researcher



PUTTING IT TOGETHER

Plan your 30-Day Pilot

A worked example showing how Pages 2–4 combine into one time-boxed pilot.

THE USE CASE

Auto-extract data from supplier invoices

A 25-person service-based SME processes about 80 supplier invoices per month. Invoice details such as date, supplier, amount and VAT are manually entered into accounting software.

FROM PAGE 2 – SCOPE

Test: Run automated invoice extraction in parallel with manual entry for one month of new invoices.

Success metric: Reduce time per invoice by at least 40% while keeping the error rate below 3%.

Stop/adjust if: Pilot isn't on track to meet success metric by day 21.

FROM PAGES 3 & 4 – DATA AND PROVIDER CHECKS

Data tier: Confidential (financial documents)

Provider/setup: CH/EU-hosted invoice extraction tool. The tool only needs to extract the required invoice fields reliably.

Oversight: AI extracts, human verifies, nothing is entered automatically.

30-Day Pilot Plan

WEEK 1

Setup

- Select pilot sample: one invoice type and new invoices received during the pilot
- Confirm provider checks: data location, data processing agreement, training use and retention
- Measure baseline: manual time per invoice and current error rate
- Configure extraction fields and human verification step

→ **Baseline confirmed**

WEEK 2

Test

- Run AI extraction in parallel
- Compare extracted fields with the manual result
- Log time and accuracy
- Keep human verification in place

→ **First results measured**

WEEK 3

Improve

- Review common errors: scans, unusual layouts or handwritten notes
- Adjust extraction rules or tool settings
- Add one invoice type only if results are stable
- Flag low-confidence or unclear fields before entry

→ **Refined workflow**

WEEK 4

Decide

- Compare time and error rate to the baseline
- Make the final pilot decision: continue, adjust or stop
- Document what worked/failed and what to fix before production

→ **Decision made to continue, adjust or stop**

Day-30 Decision

CONTINUE

Target reached: time per invoice reduced by $\geq 40\%$ and error rate $\leq 3\%$
→ Expand to other invoice types.

ADJUST

Some gain but below target
→ fix the setup, test 2 more weeks, then decide again.

STOP

Weak benefits, high error rate too or setup too complex
→ document failure, pick a new use case from Page 2.

Before scaling: common pitfalls to avoid



Skipping the baseline

No baseline number means no proof of gain. You will argue from intuition, not data.



Scaling before production is ready

A successful pilot is not production-ready by default. Check budget, contracts, monitoring, training, fallback and support before expanding.



No owner after the pilot

Without a named owner past day 30, the workflow drifts or quietly dies. Decide who owns it long-term, or stop.



Drifting past day 30

Without a hard decision date, pilots become permanent unofficial workflows by default.