WHITEPAPER

Strategic Knowledge Continuity and AI for Swiss SMEs

How companies preserve important knowledge, give AI controlled access to their data and keep people accountable.

18. Juli 2026 · 13 min read

Executive summary

Many Swiss SMEs rely on knowledge that has grown over years of daily work. It sits with experienced employees, in emails and spreadsheets, in customer histories, specialist systems and exceptions that have never been written down. Some of it is difficult or impossible for others to access.

When key people are unavailable, reduce their workload or leave the company, more than instructions disappear. The company also loses reasons, priorities and the judgement needed when information conflicts or a case falls outside the usual process. Expectations of work have changed as well. Many younger employees are less willing to accept a cumbersome process simply because it has always been done that way. They want to understand why it is necessary and have a way to improve it. This is not a criticism of earlier decisions. Established processes were often sensible under the conditions at the time. But when their purpose is no longer clear and daily work involves unnecessary searching or duplicate entry, onboarding and retention become harder.

AI can search, compare and prepare information for a decision. It becomes reliable only when it uses approved, current sources and someone is clearly responsible for checking the result. Without these foundations, its answers are difficult to assess. Because generative AI does not always produce the same answer, results and rules must be reviewed and improved in daily use.

The starting point is therefore not a model or platform, but one important workflow. For that workflow, the company organises knowledge and ownership, clarifies data access and company rules, and defines human controls. Only then does it select suitable technology and test it in a supervised trial. Because knowledge and rules remain with the company, models and tools can be changed later. AI prepares. People decide and remain accountable.

The business issue

The retirement of large age groups is increasing the shortage of skilled workers and accelerating the loss of experience. In March 2026, the Swiss SME Portal reported a KOF estimate that Switzerland could lack around 400,000 workers over the next ten years.[1] Zürcher Kantonalbank also expects retirements to exceed new entries into the labour market significantly through 2029.[2]

The challenge extends beyond retirement. Operational risk increases when important workflows function only because individual people:

- know customer exceptions and earlier commitments;
- know which source is authoritative when systems disagree;
- manage informal steps between ERP, email, spreadsheets and specialist systems;
- recognise exceptions and escalate them correctly;

- hold technical and organisational dependencies together from memory.

The result is slower onboarding, more questions and avoidable errors. Experienced employees must keep operations running while passing on their knowledge. In a labour market where long-term retention is less certain, this effort is repeated more often. Every departure creates recruitment and onboarding costs, and more context is lost. These costs and the risk to operational continuity often remain invisible in process calculations.

Knowledge continuity is therefore more than documentation. The required knowledge must be available where people work and make decisions. At the same time, responsibility must remain clear. The solution must therefore begin with the workflow, not the technology.

The starting point: a defined workflow

A specific workflow is the best place to start. It shows which knowledge is missing, which systems are involved and which decisions need protection. Models and platforms are easier to compare. But while the workflow remains unclear, the company cannot judge which technology actually fits.

First, examine one real workflow:

01 Why does it matter?

02 Which decisions and exceptions does it contain?

03 Which task, process and strategic knowledge does it require?

04 Which sources are authoritative?

05 Which data must AI access, and may it only read or also trigger action?

06 Who carries business, technical and security accountability?

07 What requirements apply to availability, freshness, traceability and cost?

08 What may AI prepare, and what remains a human decision?

The answers lead to three further questions: How should the required knowledge be organised? How can AI reach the right data safely? Which actions must remain under human control? Technology can be selected sensibly only after these questions have been answered.

Organising knowledge: three domains, three owners

Organisation starts with responsibility. When task, process and strategic knowledge are mixed in one place, it is often unclear who may approve a change. The three areas therefore remain separate. Each has a clear scope, a responsible person and an appropriate review cycle.

Three areas of knowledge, each with a clear owner

Task knowledge

Task knowledge supports a concrete action. It includes instructions, examples, case information, checklists and immediate exceptions. It is owned by the people or teams performing the work.

Process knowledge

Process knowledge connects several tasks into a workflow. It covers handovers, roles, controls, dependencies, quality rules and escalation paths. The person responsible for the overall process owns it.

Strategic knowledge

Strategic knowledge sets direction and boundaries. It includes priorities, customer and market choices, risk appetite, commitments and investment principles. It is owned by leadership or a formally accountable strategic role.

Strategy sets direction and boundaries. Processes turn them into repeatable workflows. Tasks apply those workflows to individual cases. Daily work may show that a process or strategic assumption should change. The insight is passed on as a proposal, but not adopted automatically. Task knowledge becomes process knowledge only when the process owner approves it. Process knowledge becomes strategic knowledge only when the person responsible for strategy approves it.

One subject can appear in several knowledge areas. For a customer relationship, task knowledge may concern the current case, process knowledge the handling of the relationship and escalations, and strategic knowledge its importance and long-term commitments. Scope and responsibility remain separate in each area.

This separation clarifies who may review and approve knowledge. It does not yet explain how AI can safely reach the underlying information. Sources, access and company rules must therefore be examined together.

Knowledge moves to another area only after its new owner approves it

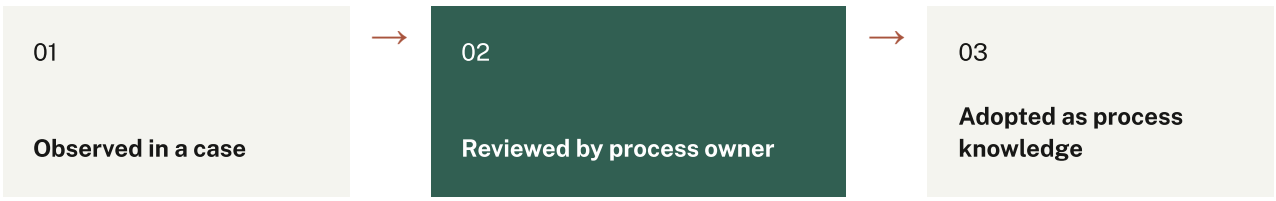

AI may suggest a change. It cannot approve or adopt it.

How employees experience the system

The knowledge architecture explains how strategy reaches concrete tasks through processes. Employees should not have to operate that architecture every day. They begin with the work in front of them:

- What needs my attention today?
- What has changed since the last review?
- What short context matters for this task?
- Which information is missing or has not been checked?
- What should happen first, and why?

A new message does not automatically become a task. AI reconciles new signals with existing work and surfaces only what needs attention because of a deadline, commitment, risk, dependency or process rule.

Priority is not an unexplained AI judgement either. The system proposes an order and states the reason. The responsible person reviews it, changes it when necessary and decides on the next action.

The employee sees a clear working picture. Behind it, approved sources, processes, responsibilities and approvals keep that picture reliable.

The daily view begins with what needs attention

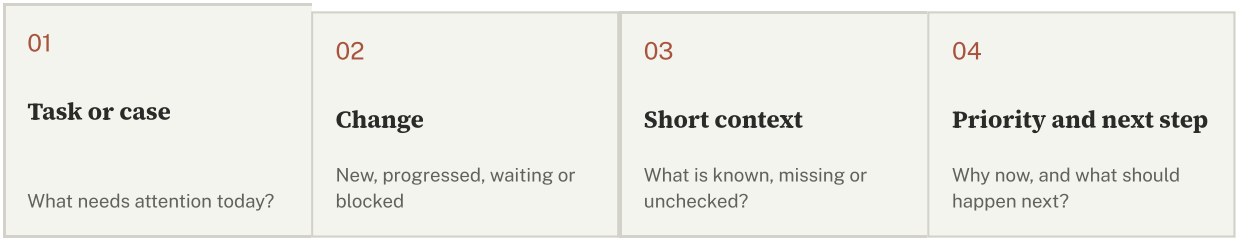

AI prepares and explains the view. The responsible person sets the order and decides.

Connecting data: clarify access and rules together

Clear responsibility alone is not enough. AI can use only information it is technically able and allowed to reach. Conversations and documents show what knowledge exists. It becomes available in daily work only when AI can access the required systems securely.

Seven points must therefore be clarified for the selected workflow:

01 **Sources:** documents, email, ERP, CRM, file stores, tickets, databases, specialist systems and human experience.

02 **Authoritative source:** which source prevails when information conflicts?

03 **Existing access:** can the systems provide the required information through search, standard interfaces, exports or controlled handovers? Reading and acting must be clearly separated.

04 **Missing interfaces:** which connection is missing, and is building it justified for this use case?

05 **Responsible people:** who is responsible for the business outcome, workflow, data, systems, security and operation?

06 **Company rules:** what is allowed for models, hosting, personal and customer data, retention and actions in other systems?

07 **Approval and traceability:** what may be displayed, drafted, stored, adopted or executed? What must remain understandable later for management, an audit or an incident review?

After this clarification, leadership can decide whether the use case is technically and organisationally feasible. It can see which data AI needs, how access could work and who must agree. It also becomes clear where data remains, what is copied, which dependencies arise and what happens when data is stale or a system fails.

Sources and access are still not enough. The company must also decide what AI may do with the information. This determines how the workflow is operated and controlled.

Operating AI: prepare work, protect decisions

When AI only reads and prepares, an error remains a proposal that a person can check. The first use therefore begins with read access. AI searches approved sources, compares cases, highlights contradictions, prepares a summary or draft, and names its sources. The result does not yet change a production system.

Once an output affects customers, money, personnel, rights or production systems, an error can have direct consequences. People therefore remain accountable for interpretation, communication, commitments and changes. Customer communication, financial obligations, personnel and legal decisions, orders, permissions and changes to critical systems require explicit approval.

Generative AI does not always produce the same answer. Sources, rules and exceptions also change. A one-time acceptance is therefore not enough. Trust develops in operation and must be confirmed repeatedly through three recurring activities:

01 **Review:** inspect results, sources, uncertainty and recurring errors.

02 **Correct:** correct the output, source, instruction, access rule or workflow.

03 **Adopt:** incorporate only confirmed improvements into the responsible knowledge area. When knowledge moves to another area, its owner decides.

This makes quality control and learning part of daily work. Five rules follow for the first use:

- a named business owner for the workflow;
- traceable sources and visible freshness;
- only the data and actions required for the use case;
- explicit approval for consequential action;
- regular review, correction and adoption.

Once the workflow, responsibility for knowledge, data access and human control are clear, implementation can proceed in stages.

Trust grows through regular review, not a one-off approval

01	02	03
Review	**Correct**	**Adopt**
Check output, source and uncertainty	Fix source, rule or workflow	Make only confirmed improvements part of the workflow

Implementation: four controlled stages

The four stages build on one another. First, they reveal dependence on individuals. Then they close knowledge gaps, test suitability in practice and define the rules for ongoing operation.

1. Knowledge and continuity check

First select one important workflow that currently depends strongly on individual people. A clear boundary keeps the effort manageable. At the same time, it reveals knowledge, sources, data flows, access, responsible people, rules and possible points of failure. The company can then decide whether to stop, address prerequisites or prepare a supervised trial.

2. Expert knowledge capture

Conversations, real cases, documents and observation make experience held by individuals available for others to review. A complete manual would be expensive and quickly become outdated. Only the knowledge required by the selected workflow is captured. Gaps and conflicts remain visible.

3. Supervised AI trial

In the trial, a clearly bounded workflow uses only approved sources and shows what each result is based on. It waits for human approval before consequential action. This keeps access, failures and approvals traceable. The trial shows whether the data and working method are suitable. Only then can possible financial value be assessed.

4. Ongoing improvement

If the trial proves useful, the workflow becomes part of normal operations. Responsibility, review cycles, access checks, monitoring, changes and incident handling are then defined as binding practices. Because business knowledge remains separate from the model and provider, the technology can change later.

The right starting workflow depends on the company. Good candidates require information from several sources to be prepared while the decision clearly remains with a responsible person.

Suitable first use cases

Handover: capture standard steps, exceptions, the order of authoritative sources and escalation criteria for one service, production or administrative workflow. The responsible expert reviews the result before handover.

Preparation for customer meetings: combine approved customer history, open commitments, previous decisions and current status into a briefing with sources and open questions. The account owner decides what is relevant and what is communicated.

Back-office preparation: for recurring enquiries, order clarification or invoice checks, AI gathers information, identifies missing data and prepares a checklist or response. The responsible person reviews the case and approves external action.

Learning across knowledge areas: confirmed observations from individual cases are proposed to the process owner. Recurring patterns become process knowledge only after approval. Leadership considers possible strategic implications separately.

The examples differ in subject matter but follow the same pattern. AI brings approved information together. A person checks the result and remains accountable for the action. Before a trial begins, leadership defines its boundaries.

What leadership must decide

Leadership decides not only which use case to pursue. It also defines the conditions under which it may operate. Knowledge, data access, workflow and approvals should not be tied to a single provider. Otherwise, changing technology later becomes expensive and risky.

Key decisions include:

- **Data and confidentiality:** which information may be processed or copied, where does it remain and when is it deleted?
- **Accountability and access:** who may view, prepare, approve and execute? Who decides when results are uncertain or wrong?
- **Dependencies:** which systems and providers are required, and what happens when data is missing or a service fails?
- **Traceability:** can sources, rules, human approvals, actions and failures be linked to a business case?
- **Cost and scale:** what does the company pay for licences, connections, operation, usage and human review? How do these costs change with wider adoption?
- **Changing or stopping:** can the company change providers or models, roll back the workflow or stop it safely without losing knowledge and control?

The use case and company policy determine the technology. By starting with a small, understandable solution, the company can test data access, quality and accountability before investing further. This does not require an enterprise-wide AI platform.

The preceding sections describe the model. I apply its principles in my own working system.

What my own working system demonstrates

The strategic vault contains long-term decisions and context. The work vault contains tasks, processes, approvals and working rules that span several projects. Subject knowledge and technical decisions remain in the relevant project workspaces.

These stores do not map the three knowledge areas one-to-one. They do, however, separate strategic direction, coordination across projects, and project-specific task and process knowledge. This keeps short-lived work out of the strategic vault and details close to their source.

AI searches, compares, structures and prepares. I decide which knowledge is adopted for the long term and which actions are carried out outside the system. I deliberately move lasting insights to the responsible store. Readable files and version history keep this path traceable. Tools and models remain replaceable.

Because I am responsible for every area, this system does not yet demonstrate operation with several owners. It does show that clear source precedence, controlled adoption, versioning, approval and replaceable tools work together.

It would be wrong to transfer this system unchanged to a company. A first engagement tests which principles fit one real workflow in that company.

A suitable first engagement

A clearly bounded knowledge and continuity check is a sensible first engagement. It provides:

- a concise description of the workflow and its failure points;
- an overview of task, process and strategic knowledge;
- named owners and approval gates;
- an understandable overview of sources, data flows, systems, access, responsibilities and controls;
- the requirements of privacy, security and governance;
- a recommendation to stop, resolve prerequisites or design a supervised trial.

Before the workflow has been examined, a benefit claim would not be credible. The engagement therefore provides a basis for decision, not a promise of success: Is the case suitable? What is missing? Is a supervised trial justified? The company does not need to select a model provider or replace a core system first.

Conclusion

Knowledge continuity exists when important knowledge is available in the workflow and responsibility remains clear. A company does not need to begin with a large platform. It can select one important workflow, clarify sources and responsibility, and test AI under human control. This turns an idea into a sound decision without tying the company to a provider too early.

Sources

[1] State Secretariat for Economic Affairs SECO / Swiss SME Portal, “Increasing labor shortage in Switzerland”, 25 March 2026: https://www.kmu.admin.ch/kmu/en/home/new/news/2026/increasing-labor-shortage-switzerland.html

[2] Zürcher Kantonalbank, “Die Boomer-Welle rollt mit grosser Wucht an”, 2025: https://www.zkb.ch/de/blog/kmu/unternehmensnachfolge-interview-boomerwelle.html