contovista

Use Case Booklet

Agentic Workflows in Lending

How Swiss banks increase efficiency and client engagement in lending.

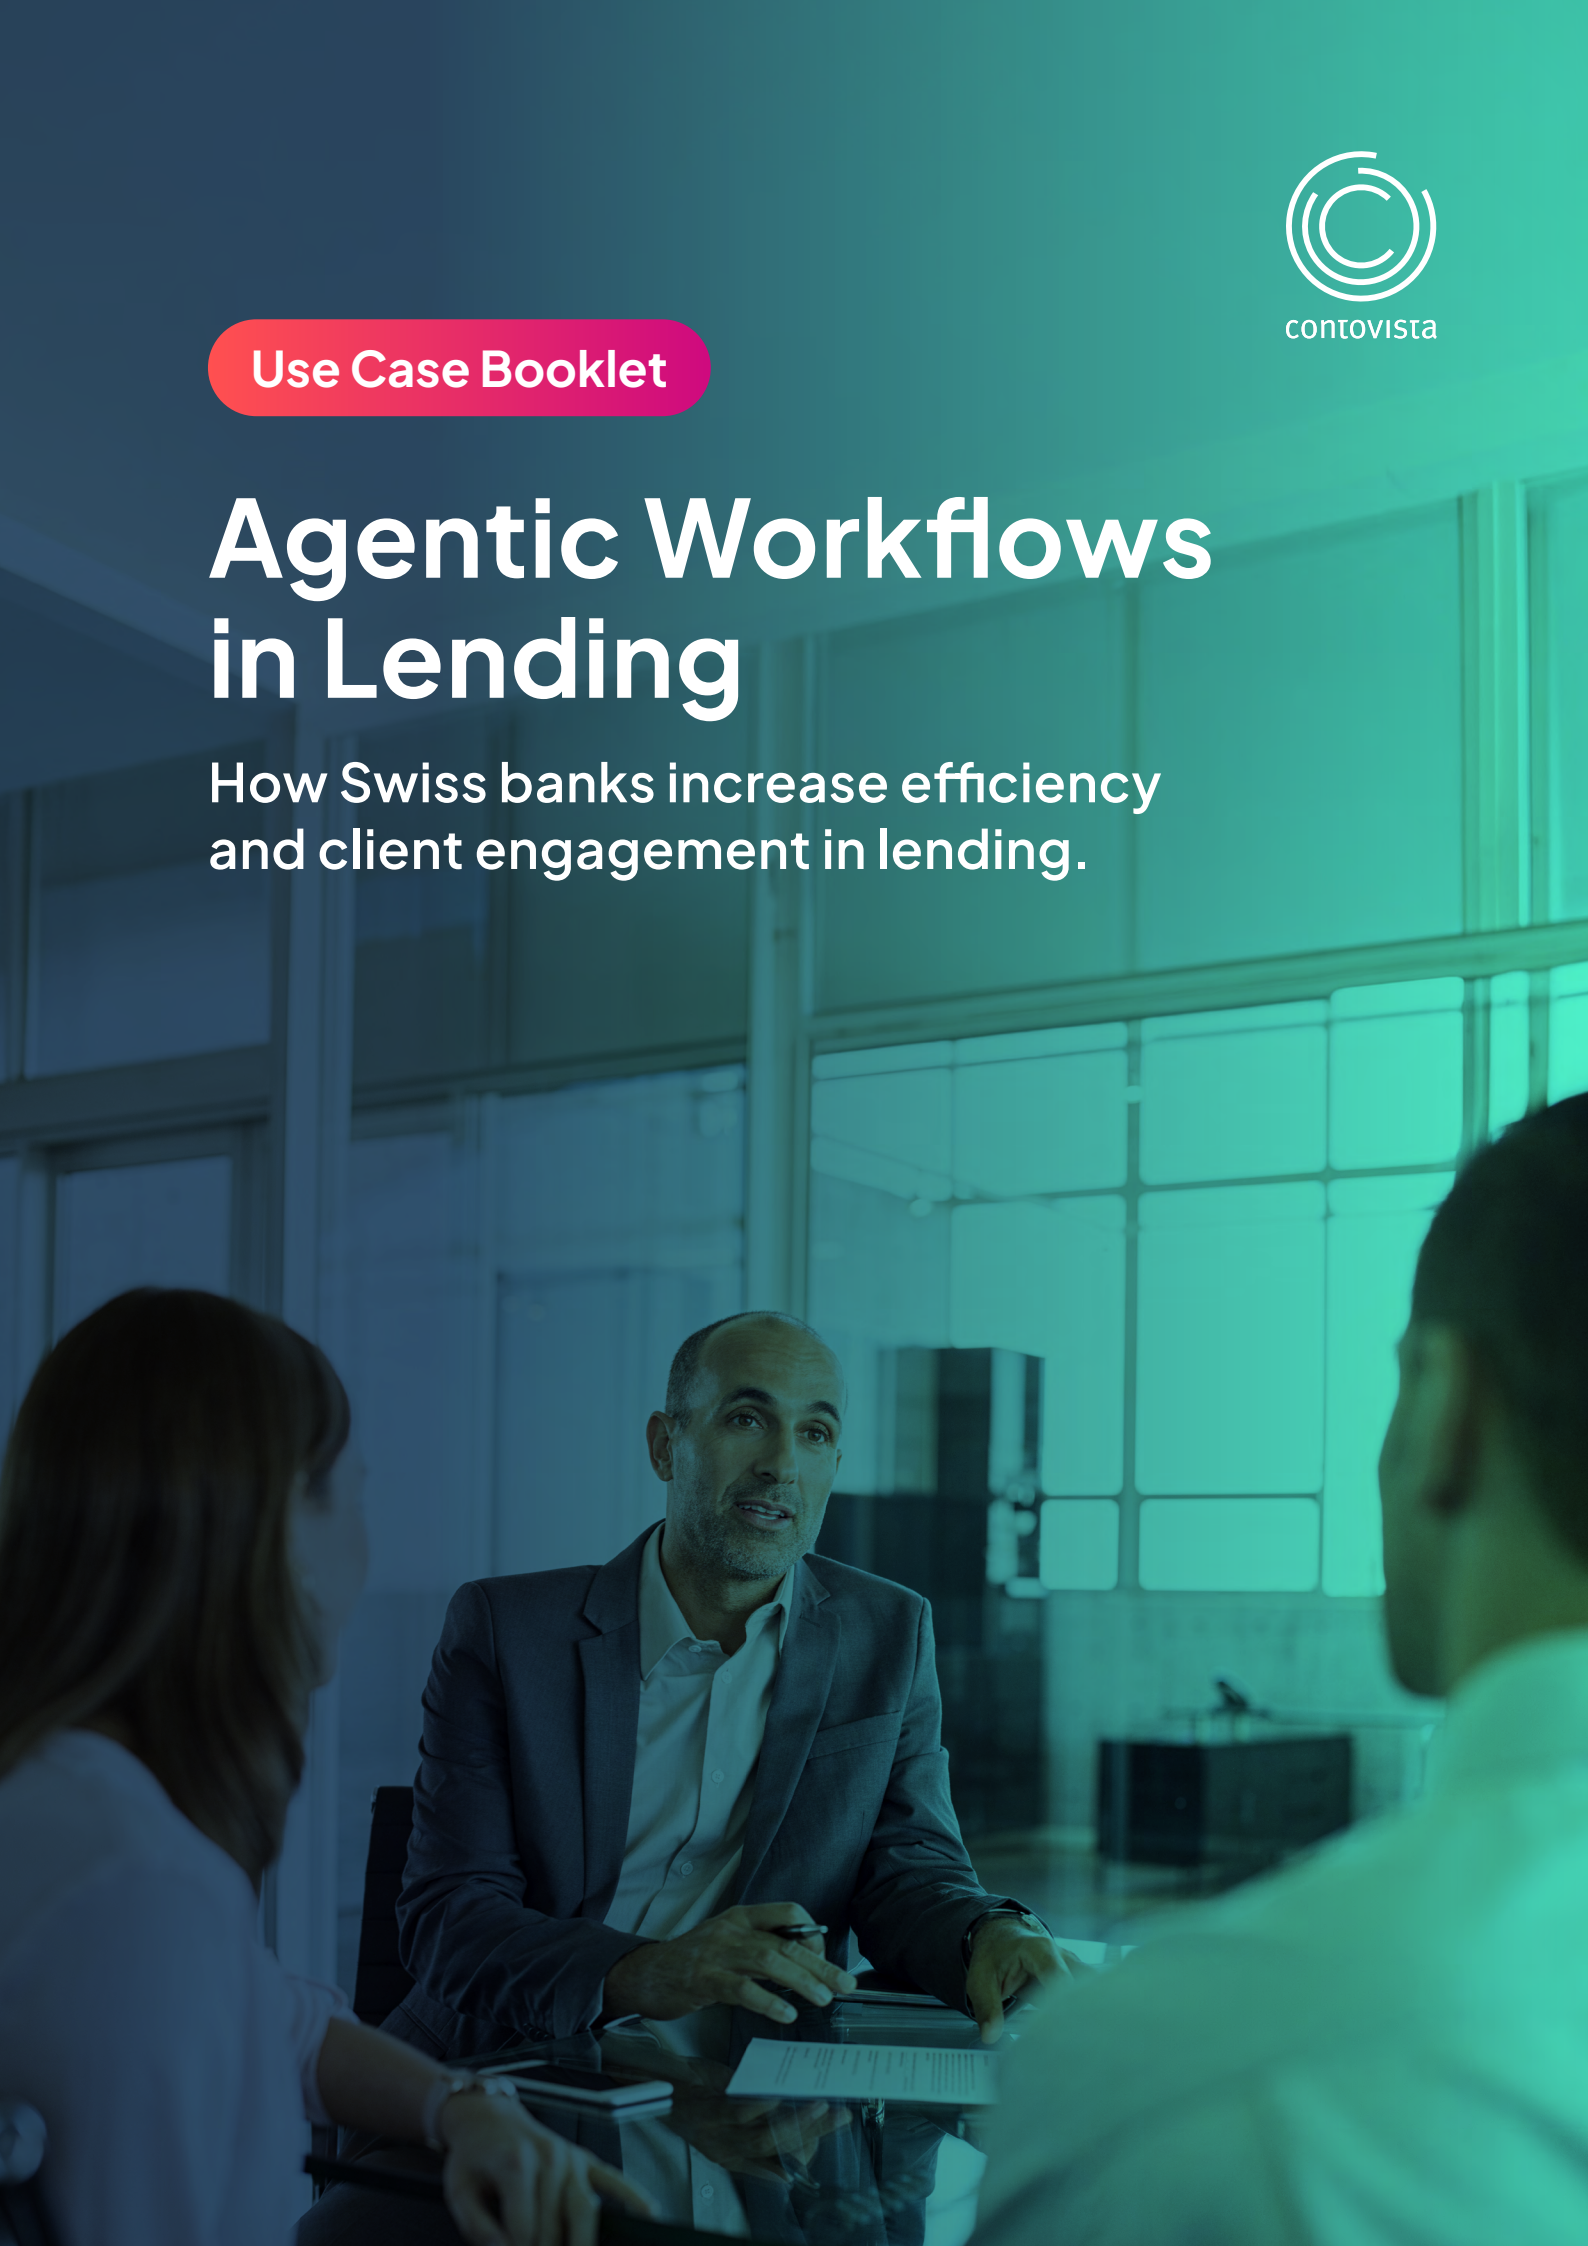

Lending today: where time, quality and transparency are lost

Lending processes at Swiss banks are complex – and not because the underlying decisions need to be. The real problem lies in how information is gathered, prepared and passed on: manually, in fragments, dependent on individuals.

Relationship managers spend a considerable share of their time on coordination rather than advice. Credit Operations struggle with incomplete dossiers and repeated requests for missing documents. Risk and Compliance wait for structured decision papers that are rarely complete on the first attempt.

The result is turnaround times that bear no relation to the actual complexity of many cases. Consistency suffers when different teams assess the same application against different standards. The rework rate is high. And the cost per case rises while margins remain under pressure.

For clients, this manifests in direct ways: they submit the same documents several times. They receive no proactive information about their financing scope. They wait for decisions and often learn of a rejection weeks later than necessary.

The Accenture Global Banking Consumer Study, a survey of more than 49,000 banking clients across 39 countries, shows that **speed and transparency in the lending process are among the strongest drivers of client satisfaction and retention.** Banks that perform well here demonstrably achieve higher growth rates.

This is not a new insight. What is new is that the technology to change it is now ready for practice.

Target picture: faster, more consistent, controlled

Agentic AI supports the lending process precisely where efficiency and quality are lost today, without replacing human control over lending decisions.

The concrete promise:

- Relevant client data is automatically consolidated and checked for completeness
- Decision bases are prepared in structured form before they reach a credit officer
- Exceptions and escalations are cleanly identified and correctly routed
- Approvals, judgements and the client relationship remain with people

The result: faster decisions, less rework and full control, combined with a better client experience.

The guiding principle is called Human-in-the-Loop – not as a constraint, but as a design decision. Agentic AI takes over repetitive, data-driven process steps. Bank staff are responsible for judgement, exceptions and relationship. It is this combination that secures ROI, compliance and credibility at the same time.

What is Agentic AI – and what sets it apart from previous AI approaches?

Agentic AI refers to systems that independently orchestrate multi-stage workflows. Whereas a classic AI system responds on request, an AI agent plans and coordinates sequences of actions across several steps, drawing autonomously on data sources, decision logic and external tools.

For the lending process, this means: an agent analyses an incoming financing enquiry, assesses creditworthiness on the basis of transaction data, identifies missing documents, produces a decision proposal and forwards it to the right place – without manual intermediate steps.

The economic difference lies in the process impact. Whereas simpler AI tools act primarily on individual tasks, Agentic AI addresses end-to-end workflows. This is precisely where the greatest gains in efficiency and quality arise.

The Contovista approach draws a clear distinction between two layers: specialised ML models deliver precise signals from transaction data such as categorisation, pattern recognition and risk scoring. Lean LLMs translate these signals into workflow orchestration. The combination lowers costs, increases traceability and enables controlled scaling.

The four use cases at a glance

	Use Case	Benefit dimension
1	Proactive pre-qualification	Client experience / Innovation
2	Intelligent application preparation	Efficiency
3	Decision preparation and intelligent routing	Efficiency / Risk
4	Decision communication and client follow-up	Client experience

The four use cases map the lending process from front to back. They can be introduced individually or operated as an integrated system, depending on the bank's priorities and integration maturity.

Use Case 1: Proactive pre-qualification

Benefit dimension: Client experience / Innovation

Problem and outcome at a glance

Transaction data, income trends and saving patterns are available, but the bank does not use them to approach clients at the right moment. Instead, the client has to take the initiative and often waits weeks before knowing what scope they can plan within.

According to the Accenture Global Banking Consumer Study 2025 (49,000 banking clients, 39 countries), proactive, personalised guidance – «Remember me» and «Delight me» – are measurable drivers of client engagement and revenue growth.

→ **Outcome:** the AI agent identifies life events and behavioural patterns from transaction data, calculates a realistic financing scope and triggers a context-sensitive client interaction before the client takes the initiative.

contovista

How it works in day-to-day banking

Step 1 Signal detection (agent, fully automated)

The agent analyses transaction data for patterns indicating an upcoming financing need: rising saving behaviour, income increases, the ending of commitments.

Step 2 Qualification calculation (agent, fully automated)

On the basis of transaction and income data, the agent calculates a financing scope in line with the bank's internal criteria – not a binding commitment, but a reliable point of orientation.

Step 3 Approval to trigger outreach (Human-in-the-Loop: RM)

The relationship manager receives the prompt with the calculated figure and underlying signals. They decide on the type and timing of contact.

Step 4 Personalised client approach (agent, semi-automated)

The agent produces a communication draft (email, push or advisory invitation), which the RM adjusts and sends.

Step 5 Transition into the lending process (Use Case 2)

If the client signals interest, the prepared data is handed over seamlessly.

Roles and control points

Role	Task	Control point
AI agent	Signal detection, calculation, communication draft	None (fully automated)
Relationship Manager	Reviewing the signal, approving the outreach trigger, adjusting if needed	Mandatory approval before contact
Client	Responding to the approach, initiating the advisory conversation	--

Metrics for assessing the benefit

KPI	Description
Conversion rate, financing prospects	Share of approached clients who arrange an advisory conversation
Time-to-First-Contact	Time between a detected signal and the first client approach
Proactively detected occasions	Share of financing enquiries initiated through proactive outreach
NPS, financing advice	Client satisfaction in the entry process; comparison before/after implementation
Cross-sell rate	Share of additional products in the context of an initiated financing

Context: operations and regulation

The agent evaluates exclusively existing transaction data and makes no lending decisions. In data protection terms, the Swiss DSG requires transparency about the purpose and scope of data use. The Contovista platform is aligned accordingly and carries the swiss made software + AI quality label.

The mandatory approval by the RM ensures that no fully automated client approach takes place and that regulatory requirements for advisory quality are met.

Use Case 2: Intelligent application preparation

Benefit dimension: Efficiency

Problem and outcome at a glance

Clients upload documents, fill in forms manually and frequently still fail to provide all the required papers on the first attempt. Credit Operations spends a disproportionate amount of time on completeness checks and follow-up rounds before the actual credit assessment can even begin.

→ **Outcome:** the AI agent extracts data automatically from uploaded documents, populates the credit application, checks completeness and delivers Credit Operations a structured, decision-ready dossier – without manual data entry.

contovista

How it works in day-to-day banking

Step 1 Document intake and extraction (agent, fully automated)

The agent recognises document types (salary statements, tax returns, account statements), extracts all relevant fields and transfers the data into the application in structured form. Illegible documents are flagged for manual processing.

Step 2 Completeness check (agent, fully automated)

The agent compares the intake against the documents required for the respective credit type. Missing or contradictory information is identified immediately.

Step 3 Automated follow-up request (agent, semi-automated)

For incomplete dossiers, the agent produces a precise follow-up request to the client – specifying the missing documents and submission options.

Step 4 Consistency check (agent, fully automated)

The agent checks whether stated income data is plausible against the transaction history and flags anomalies for human review.

Step 5 Handover to Credit Operations (Human-in-the-Loop: credit officer)

Credit Operations receives a complete, structured dossier and begins directly with the substantive review.

Roles and control points

Role	Task	Control point
AI agent	Extraction, completeness check, follow-up request, consistency check	Flagging of anomalies for human review
Client	Document submission, follow-up delivery if needed	--
Credit officer	Substantive credit review on the basis of the prepared dossier	Final dossier before further processing
RM (optional)	Support in complex cases	Escalation in the event of systematic gaps

Metrics for assessing the benefit

KPI	Description
Processing time per dossier	Average time from intake to a decision-ready dossier
First-Time-Complete-Rate	Share of dossiers complete on first intake
Follow-up query rate	Number of manual follow-up rounds per application
Cost per Application	Total cost of the administrative dossier preparation per case
Data-entry error rate	Share of erroneous or inconsistent fields in the application

contovista

Context: operations and regulation

The agent makes no lending decisions whatsoever; the substantive assessment remains entirely with the credit officer. The requirements of the DSG apply to the processing of client documents; automated processing is permissible where the purpose is clearly defined. All steps are logged and auditable as a basic prerequisite for FINMA-compliant processes.

Use Case 3: Decision preparation and intelligent routing

Benefit dimension: Efficiency / Risk

Problem and outcome at a glance

Creditworthiness, risk profile, ratios, alignment with credit policies: the credit officer has to consolidate all of this manually before being able to reach a decision. Added to this is the question of correct routing: poor escalation management creates unnecessary waiting times or decisions at the wrong hierarchical level.

→ **Outcome:** the AI agent calculates creditworthiness and risk assessment, produces a structured decision proposal and routes the case automatically to the right place – with a clearly defined escalation path.

contovista

How it works in day-to-day banking

Step 1 Dossier analysis and risk scoring (agent, fully automated)

The agent evaluates the prepared dossier and calculates a traceably documented risk score: income stability, debt servicing capacity, asset development, risk signals.

Step 2 Alignment with credit policies (agent, fully automated)

Standard cases are flagged as such. Deviations from defined parameters are explicitly highlighted.

Step 3 Production of the decision proposal (agent, fully automated)

The agent produces a structured credit proposal with a creditworthiness assessment, recommended terms, identified risks and line of reasoning – explicitly declared as a decision aid, not a decision.

Step 4 Intelligent routing (agent, fully automated)

The agent routes the case to the correct decision level: standard to the credit officer, complex cases to the senior level, regulatorily sensitive cases with a compliance note.

Step 5 Decision (Human-in-the-Loop: credit officer)

The credit officer takes the decision and can approve, deviate or escalate. Decision responsibility lies entirely with the human.

Roles and control points

Role	Task	Control point
AI-Agent	Risk scoring, policy alignment, decision proposal, routing	Flagging of all deviations for human assessment
Credit officer	Substantive assessment and decision	All lending decisions, without exception
Senior credit officer	Decision on complex or deviating cases	Mandatory escalation at defined thresholds
Compliance / Risk	Review of regulatorily sensitive cases	On a compliance note from the agent

Metrics for assessing the benefit

KPI	Description
Time-to-Decision	Time from a complete dossier to a lending decision
Routing accuracy	Share of correctly routed cases on the first attempt
Consistency of the risk assessment	Deviation rate between the AI proposal and the final decision
Escalation rate	Share of cases that must be decided at a higher level
Decision quality	Share of loans requiring subsequent rework or adjustment

Context: operations and regulation

FINMA requires traceability of lending decisions, clear responsibilities and documented lines of reasoning. The agent takes no decisions; it produces proposals that are fully auditable. Deviations from the proposal, as well as approvals, are documented in equal measure. In practice, this creates a better audit-trail quality than manual processes.

The basis is the Contovista Enrichment Engine with industry-leading categorisation accuracy – tailored specifically to Swiss banking conditions.

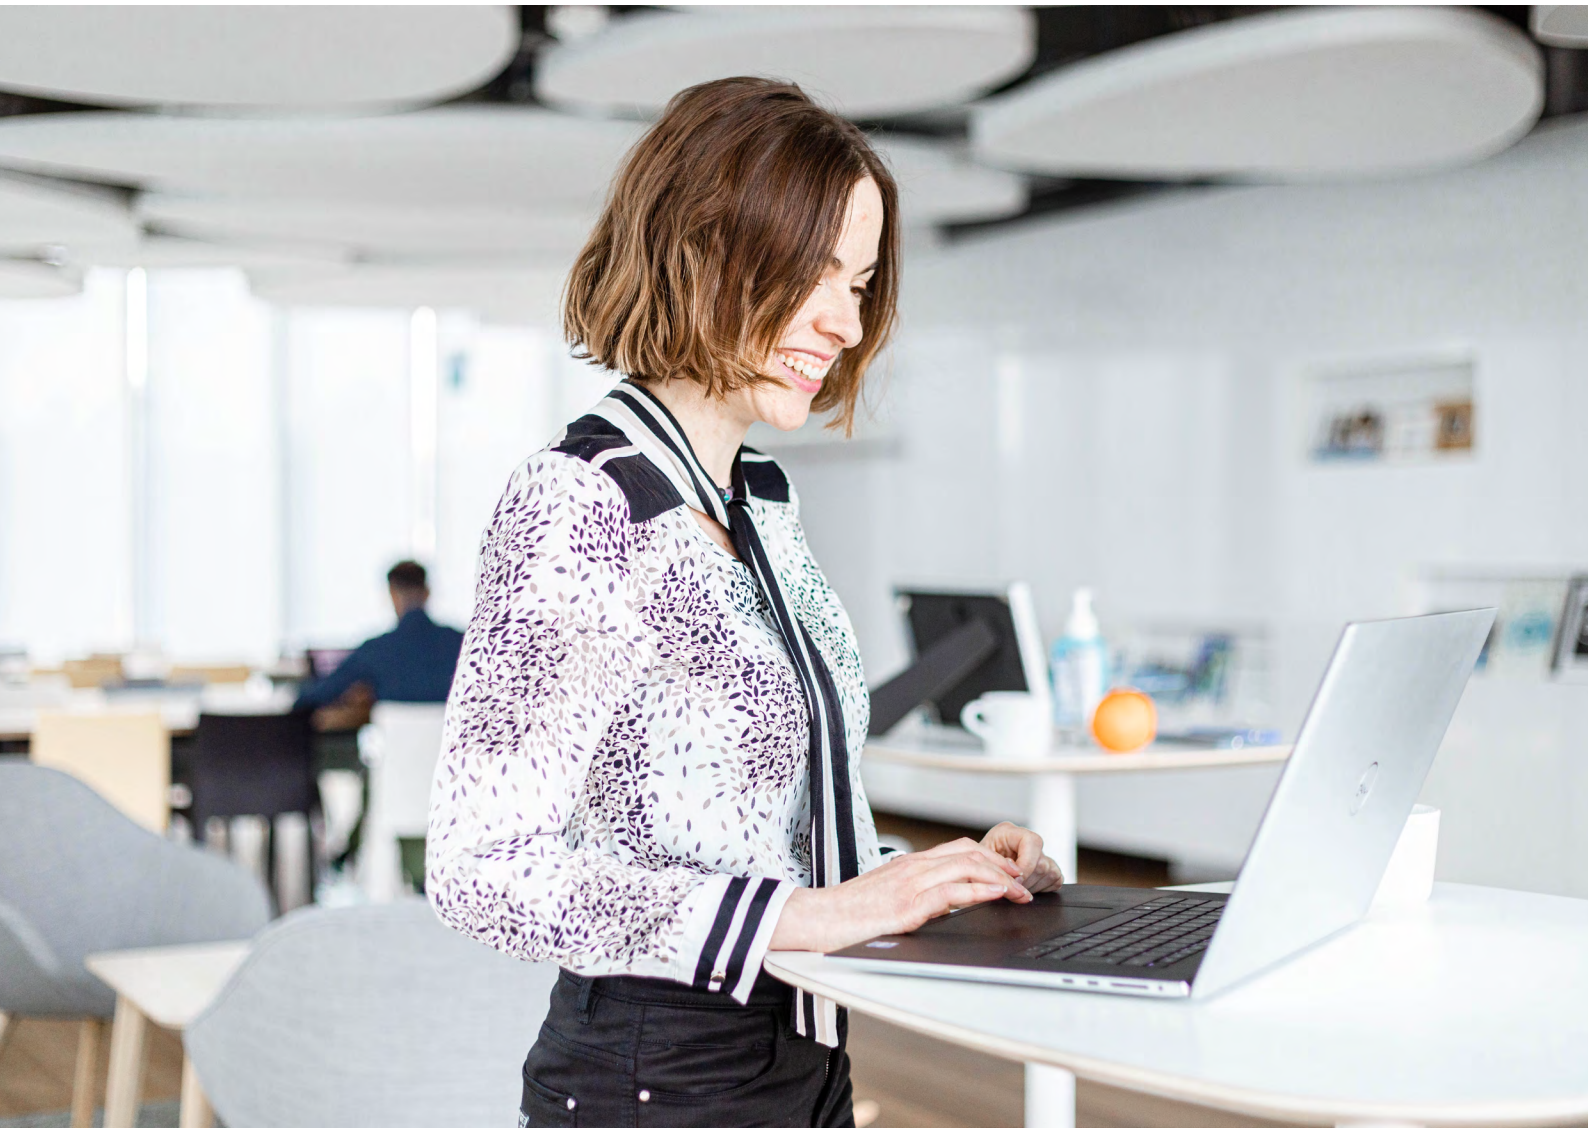

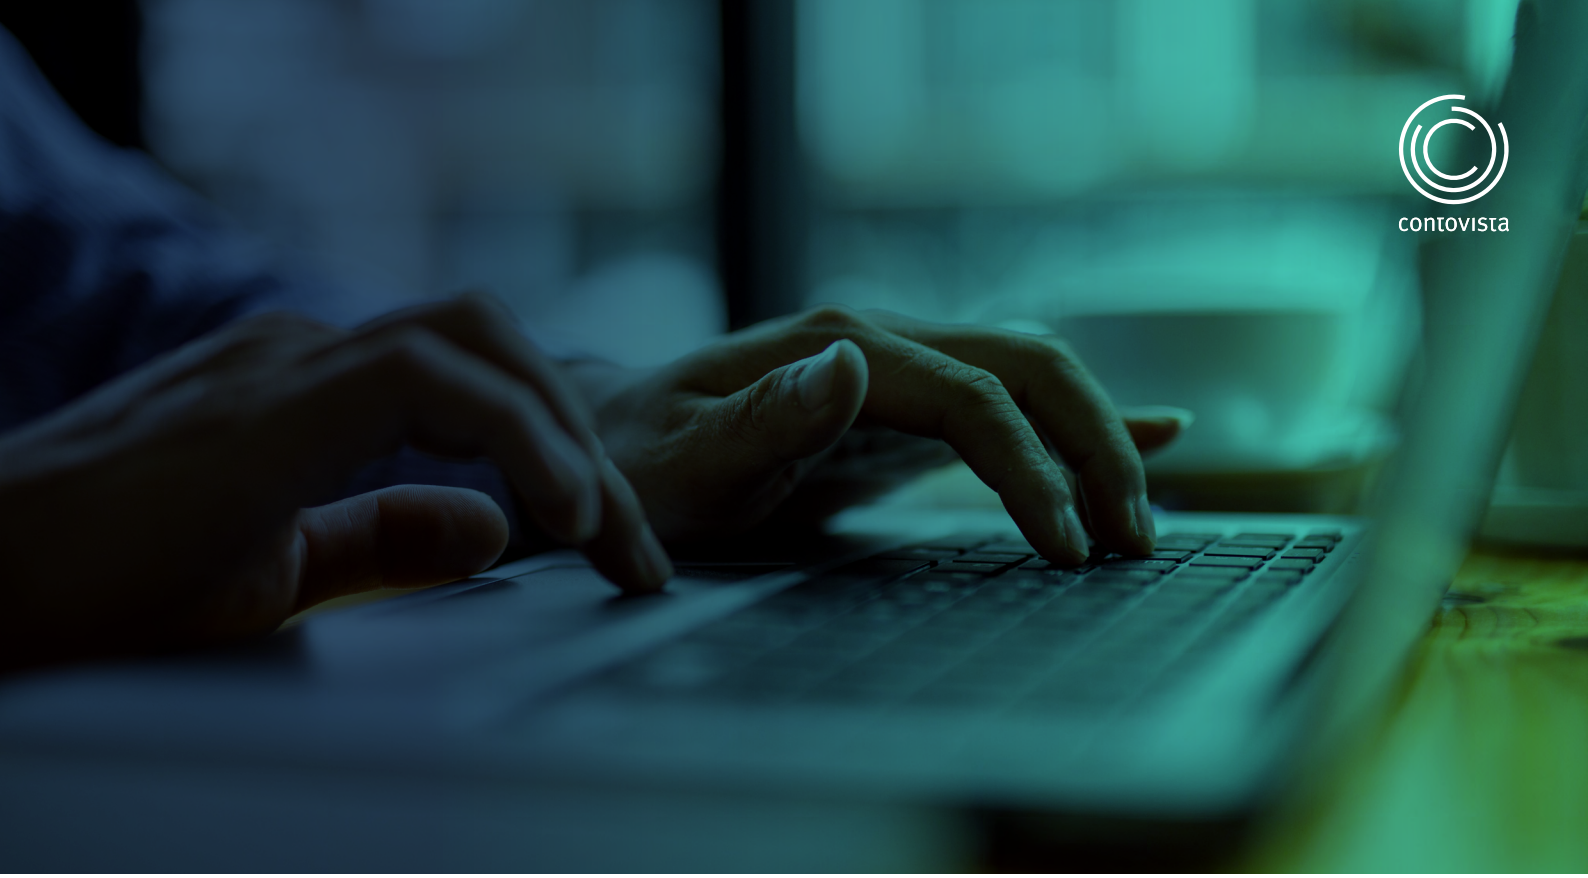

Use Case 4: Decision communication and client follow-up

Benefit dimension: Client experience

Problem and outcome at a glance

Decisions are communicated too late, rejections explained inadequately, next steps not clearly named. With negative decisions in particular, the delay is serious: clients have no basis for planning, and the bank loses the chance to strengthen the relationship even after a rejection.

→ **Outcome:** immediately after the decision, the AI agent produces structured communication – for approvals with clear next steps, for rejections with a comprehensible justification and concrete options for action.

contovista

How it works in day-to-day banking

Step 1 Context preparation (agent, fully automated)

As soon as a decision is recorded in the system, the agent pulls the full case context from Use Case 3: decision bases, terms or grounds for rejection.

Step 2 Production of the communication draft (agent, fully automated)

For approvals: confirmation, terms, next steps, point of contact. For rejections: a comprehensible justification on the basis of the documented decision bases, supplemented by concrete options for action and possible alternatives.

Step 3 Approval by the RM (Human-in-the-Loop)

The RM reviews and approves. In sensitive cases, the agent recommends a personal conversation rather than purely digital communication.

Step 4 Dispatch and follow-up (agent, fully automated)

Dispatch via the preferred channel, documentation, automatic reminder to the RM if there is no response

Step 5 Client follow-up (semi-automated)

For rejections, the agent monitors client signals in the following weeks and, if there is a change, triggers a renewed prompt to the RM (link to Use Case 1).

Roles and control points

Role	Task	Control point
AI agent	Context preparation, communication draft, dispatch, monitoring	Flagging of sensitive cases for personal contact
Relationship manager	Review and approval of the communication	✓ Mandatory approval before dispatch
Client	Receipt, response, decision on next steps	--

Metrics for assessing the benefit

KPI	Description
Time-to-Notification	Time between a decision and informing the client
NPS after a negative decision	Client satisfaction after a rejection – often more differentiating than NPS after approval
Query rate after communication	Share of clients with queries after receiving the decision communication
Re-application rate	Share of rejected clients who enquire again at a later point
RM time saved on follow-up communication	Reduction in time spent on decision communication per case

contovista

Context: operations and regulation

In Switzerland, rejections must be justified; clients have the right to understand the decision bases. Here the agent raises the quality and consistency of how justifications are presented compared with manually produced communications. The mandatory approval by the RM ensures that human assessment of appropriateness is retained, particularly in sensitive cases.

How does a bank get started?

Agentic AI in the lending process does not have to begin as a major project. The economically viable path runs through focused piloting with clearly defined processes, robust data and measurable objectives.

Contovista recommends a structured entry over four weeks:

Week 1
Prioritisation and process analysis

A joint assessment of the four use cases against three criteria: process maturity (how standardised is today's workflow?), data quality (how reliable is the available transaction data?) and business impact (where is the ROI most directly measurable?).

In most cases, Use Case 2 (application preparation) is suitable as an entry point, because it delivers quickly visible efficiency gains and touches no decision responsibility. Use Case 1 (pre-qualification) offers the strongest strategic lever for differentiation for banks that prioritise client engagement.

Week 2
First functional version

Building the technical connection to the Contovista infrastructure, configuring the workflow logic for the prioritised use case, first live tests with real data in a secured environment.

Weeks 3 and 4
Pilot operation and optimisation

Operation with a defined subset of real lending cases. Systematic evaluation of the KPIs. Adjustment of parameters and thresholds. Involvement of the affected teams (RM, Credit Ops) in quality assurance.

After the pilot
decision and scaling

On the basis of the pilot data, the bank decides which use cases to expand and in what order. The framework is modular: each of the four use cases can be operated independently or linked as an integrated system.

Experience shows that banks starting with a clearly defined use case not only achieve ROI faster. They simultaneously build the internal trust needed to extend to further workflows.

Efficiency without loss of control: Agentic Workflows in lending with Contovista

Making lending processes more efficient without giving up control – and improving the client experience at the same time: this is no contradiction. With Agentic AI and a consistent Human-in-the-Loop approach, both goals can be achieved together. Contovista delivers the necessary infrastructure for this – proven in practice, aligned with Swiss requirements and deployable in a modular way.

The entry is already worthwhile with a single use case. Those who start with clearly defined processes, robust transaction data and measurable objectives quickly see what is possible – and simultaneously build the internal trust needed for scaling. This is how a pilot becomes a lasting competitive advantage.

Do you want to learn how Contovista can support your bank – with AI strategy and beyond?

Talk to our experts now

contovista

contact us
info@contovista.com

www.contovista.com

Contovista
Wiesenstrasse 5
8952 Schlieren
+41 44 730 00 32