Sidetrade Launches SAFE, the Sidetrade Agentic Framework for Enterprise

Sovereign Agentic Framework for CFOs:

- Keep sensitive customer data away from every Frontier AI Lab and hyperscaler
- Match every agent action or prompt with the most accurate and cost-efficient fine-tuned open-weight model
- Bring the full stack required for agentic Order-to-Cash under one roof, from the proprietary Data Lake and O2C-specialized models to private data centers and owned GPU compute
- Secure, control and track every financial operation an AI agent performs (ISO 27001 and SOC 2 Type II compliance)
- Lock running cost visibility through multi-year contract commitments to deploy AI agents at scale

September 10, 2026 | Billion-dollar enterprises are being asked to put their receivables data into models they do not operate, and to sign for AI agents whose running cost is set by another company's price list. Neither is a question about model quality, and neither gets solved by a better model. Today, Sidetrade (Euronext Growth: ALBFR.PA), an AI-native company dedicated to Order-to-Cash, answers both by launching SAFE, the Sidetrade Agentic Framework for Enterprise. Every agent built on SAFE runs on Sidetrade's own models, in Sidetrade's own data centers, and on Sidetrade's own token production.

Olivier Novasque, CEO and Founder of Sidetrade, stated:

"During the first wave of generative AI, the question every enterprise asked was which model was the most powerful. Finance departments are now asking a different question: How do we put AI to work on one of our most sensitive processes, cash flow generation, without handing control of our customer data, operating model and intelligence to someone else? And how can we invest in agentic AI at scale if we have no visibility into what it will cost to run in such a volatile environment? Over the past few months, we have seen CFOs become increasingly focused on the sovereignty Sidetrade offers. Addressing those concerns takes the full stack. You have to own the infrastructure, the compute, the models and the framework the agents run on. We do, from our data centers to Aimie IQ. But owning the stack only makes AI sovereign. What makes it intelligent is the data, and that is the layer we own that no AI company can easily buy or replicate. The Sidetrade Data Lake manages nearly 10 trillion dollars in B2B transactions across close to 45 million corporate buyers, accumulated over more than a decade. SAFE is where the stack and the data meet. It is what enabled our AI-native products to win nearly every competitive deal in the first half of 2026, and it is what lets our customers know upfront what their AI will cost to run at scale. Sovereign. Intelligent. Predictable. No one else in the Order-to-Cash market can offer all three."

The Gap SAFE Addresses Is Measurable

According to McKinsey (State of Organizations 2026), **88% of organizations have deployed AI, but 81% have yet to see any significant financial impact** from it. The gap is widest in finance.

An AI agent that touches receivables, disputes and customer credit must be connected to the right data, act within policy, escalate deterministically when it should not decide alone, leave a complete audit trail, and improve from one run to the next. Most agentic deployments fail on one of these points rather than on the quality of the model.

On top of these operational challenges, large enterprises face another set of barriers to deploying agentic AI at scale: **uncertain economics**, with no visibility into token consumption or future per-million-token pricing; **data sovereignty concerns** when sensitive financial data flows through third-party models and hyperscalers; and **governance requirements**, particularly the ability to define, control and audit what an autonomous agent is authorized to do.

SAFE was built to answer all.

SAFE, the Sidetrade Agentic Framework for Enterprise

SAFE is the **common foundation under Sidetrade's three AI-native products**: **Aimie IQ,** the engagement and orchestration layer accessible in natural language; **Aimie Agents**, a library of ready-to-deploy agents, starting with the Aimie Cash Collection Agent; and Customer Agents, built by customers or their partners in **Agent Builder Studio**. The result is one AI-native framework with the same controls for every agent, regardless of who builds them.

SAFE is organized into two layers. The **Agent Connect** handles everything that surrounds an agent. It combines an **Agent Registry** and a single declarative configuration model, an Event Router that dispatches Order-to-Cash (O2C) events to the right agents, and a **Skills Registry** of reusable O2C skills that agents share and invoke. The **Agent Harness** governs the agent at runtime. It assembles context from O2C data, orchestrates reasoning on Sidetrade's models, selects the right tool among O2C actions, enforces policy with deterministic escalation, and maintains audit trails and memory so that each run improves the next.

Enterprise readiness is built into SAFE rather than added afterward. Every agent action is logged and auditable end to end. Policy controls and guardrails define what an agent can execute autonomously, when human intervention is required, and who receives the escalation. Before an agent can be published to the Agent Registry, it is validated against Sidetrade's security standards and best practices.

Luke Hennerley, VP of AI Operations at Sidetrade, stated**:** "The last era of enterprise technology was about accumulating data. The AI-native era is about accumulating learning and turning it into intelligence. Sidetrade has both. We start from a proprietary Data Lake that gives our AI context and history at a depth no incumbent can match. SAFE is how we put it to work. We run everything on our own GPUs, so the inference is ours, and now the learning is too. Memory, playbooks, feedback, what worked and what did not, none of it leaves our four walls. When a decision needs a human, Aimie remembers what they chose and why, and it is there for the next time. Our users and agents working together, learning side by side in the SAFE environment, creates the ultimate learning loop. That is how AI investment compounds with usage."

Sidetrade Brings the Full Stack for Order-to-Cash Under One Roof

Running an enterprise AI agent takes five layers. Almost no company controls all five.

Layer	Frontier AI Labs	Hyperscalers	Enterprise Application Vendors	Sidetrade
Cloud and data centers	Rented or co-invested	Owned, public	Rented	Private
GPU compute	Rented or co-invested	Owned, shared	Rented	Owned
Fine-tuned open-weight models	Owned, closed weight	Resold	Rented	Private Customized
Agent framework and harness	Owned	Partial	Partial	Owned
Cross-network domain data	None	None	Own customers only	Owned

Typical architectures based on public disclosures

Sidetrade operates interconnected, secured data centers on both sides of the Atlantic, run by its own teams, inside its own private cloud. It runs its own GPU compute and produces its own tokens. Its models, a selection of large language models (LLMs), are fine-tuned for Order-to-Cash on open-weight foundations, alongside proprietary predictive models for payment behavior, buyer risk and next best collection action.

Owning the full stack is what makes Sidetrade's AI sovereign. Sensitive financial data stays inside the customer's boundaries and is never used to train any third-party models. Models are controlled and customized by Sidetrade. Compute is private and its cost is known and guaranteed to its customers, enabling them to make an agentic investment decision for the medium to long term. Agents in production are fully auditable.

The model layer is Sidetrade's to choose, so neither Sidetrade nor its customers are locked into a single model vendor's roadmap, pricing or availability. As open-weight models advance, SAFE can adopt them with no change to the customer's data exposure or cost.

Sidetrade's AI infrastructure is certified ISO 27001 and SOC 2 Type II; agents are validated against OWASP's Agentic Top 10 before publication to the Agent Registry. For a finance department, that answers the question a receivables agent raises first. An agent working disputes, credit and collections handles commercial and financial data a company holds. On SAFE, that data never leaves the perimeter it started in.

The Data Layer That Makes the Difference Behind SAFE

A frontier model, however powerful, knows nothing about how a given buyer pays, which disputes tend to resolve, or which collection strategy works for which customer. Sidetrade has spent more than ten years accumulating that knowledge in its proprietary O2C Data Lake inside its private cloud. Today it aggregates close to $10 trillion in B2B transactions across nearly 45 million buying companies. Every model in the stack is trained on it, and **no model provider, hyperscaler or competitor can buy or build it in the short or medium term**. "This choice is ten years old, not ten months old," said **Novasque**. "Sidetrade began building its O2C data asset in 2015 and moved its models onto its own GPUs before the agentic wave made compute scarce. SAFE completes that stack."

Sidetrade has built a living system rather than a static asset. With SAFE, every agent run, every sleep cycle, every skill created in Agent Builder Studio and every piece of user feedback flows back into a shared body of O2C intelligence: skills, agents, memory and knowledge.

The Sidetrade Data Lake grows with every customer onboarded on the O2C intelligence platform, and every customer benefits from the intelligence of all the others while their own data stays isolated. This is the **Sidetrade Network** effect.

Sovereignty by Design: The Right Model for Every Task, at a Predictable Cost

Sovereignty in SAFE is not a contractual promise. It is a consequence of how the framework works. Not every task in O2C calls for an LLM, and few call for a frontier one. Predicting a payment date or scoring customer risk is the work of Sidetrade's proprietary predictive models. Routing an event, applying a policy or executing a workflow step is deterministic code. Only the tasks that need language, such as reasoning over a complex dispute or drafting a message in a customer's tone, reach the O2C fine-tuned LLM. SAFE directs each task, by use case, to the most appropriate model in the stack, and every one of those models runs on Sidetrade's own compute. **No prompt, no document and no customer record is ever sent to a public model provider or a hyperscaler**, because the stack has no need for it.

The same design protects the cost. Agentic AI consumes far more tokens than a chat interface, because agents reason, call tools, read results and reason again. Across the market, this is pushing vendors from subscriptions toward metered, consumption-based pricing, and pushing finance departments to hesitate before deploying agents at scale. **SAFE spends tokens only where they create value**. Context assembly gives each model call exactly the O2C data it needs. Routing keeps simple tasks away from expensive reasoning. Deterministic escalation stops an agent spending tokens on a decision that policy reserves for a human. And because skills and memory accumulate with every run, agents do not rediscover the same knowledge twice. For a vendor dependent on a third-party API, the cost of a task is fixed by someone else's price list. **On SAFE, the cost falls as the system learns.**

This is why Aimie IQ and the agents that run on SAFE are sold on subscription**, including a volume of tasks and "O2C IQ" (token equivalents)**, with no exposure to third-party per-million-token (MTok) pricing. Sidetrade's customers keep control of their data and gain lasting visibility over the cost of their agents in production.

Olivier Novasque, CEO and Founder of Sidetrade, added: "How can a CFO sign off on a significant investment in AI agents without knowing what the running costs could be in six, twelve or twenty-four months, especially when per-million-token prices are so volatile? How can they calculate their ROI? At Sidetrade, we decided to guarantee compute pricing for our customers through multi-year contracts, so their agentic investments are both predictable and sustainable. That is what sovereignty means for the Office of the CFO. You control the data, you control the models, and in the end you control the compute and the bill."

SAFE is already in production under the Aimie Cash Collection Agent, deployed with multinationals including Securitas, Accor and Sodexo. It will underpin the production rollout of Aimie IQ and Agent Builder Studio to beta customers in the fourth quarter of 2026, in line with the timetable announced on July 21, 2026.

Sidetrade applies the same framework internally so that intelligence learned in production and intelligence built inside the company compound in one place. "We ship the framework we run on," said **Mark Sheldon, Chief Technology Officer of Sidetrade.** "Our own teams build agents on SAFE every day, on the same stack, with the same policy controls and the same observability our customers get. With this inside-first approach and in combination with AI-driven development, in my ten years at Sidetrade I have never seen a project move so fast. We see a path where we expect our platform to be getting more intelligent daily and to reach the stage in the near future, with loop engineering and eval-first skill creation, where it will effectively be building itself. The future is incredibly exciting."

The launch of SAFE marks a further step in the rollout of the O2C Intelligence 2030 plan, which targets more than 50% of revenue from AI-native products and an EBITDA margin of 30% to 35% by 2030.

Next financial announcement

First Half Year Results for 2026: September 22, 2026 (after stock market close)

Investor relations @Sidetrade

Christelle Dhrif +33 6 10 46 72 00 cdhrif@sidetrade.com

Media relations @Sidetrade

Oli Thornton +44 7933 108 107 oli.thornton@sidetrade.com

About Sidetrade (www.sidetrade.com)

Sidetrade (Euronext Growth: ALBFR.PA) is an AI-native company dedicated to Order-to-Cash (O2C). Its platform combines O2C applications, autonomous AI agents, and Aimie IQ, an intelligent natural-language interface that helps multinationals accelerate their cash generation, all built on SAFE, the Sidetrade Agentic Framework for Enterprise. Sidetrade operates the world's largest proprietary O2C Data Lake, built on nearly $10 trillion in B2B transactions and close to 45 million buying companies. This data is used to train specialized AI models that monitor, analyze, decide, and act autonomously throughout the O2C cycle. Sidetrade supports companies in 85 countries with 450 employees across Europe, North America, and Asia-Pacific.

For more information, visit us at www.sidetrade.com and follow us on LinkedIn at @Sidetrade.

In the event of any discrepancy between the French and English versions of this press release, only the English version is to be taken into account.