

Serving Machine Customers

Considerations for Financial Institutions

Institute of Financial Services Zug IFZ
www.hslu.ch/ifz

finnova.

 **FINSTAR**

 **inventxLAB**
INNOVATION BY DESIGN

 Kanton Zug

 **SFTI**

 **SIX**

 **Zürcher
Kantonalbank**



Contents

1	Introduction	2
2	Social Dimension	3
3	Technological Dimension	5
3.1	Digital Experience Platform	5
3.2	Decoupling Platform	6
3.3	Business Platform	7
3.4	Hybrid-Cloud Platform	7
3.5	Data Factory	8
3.6	Security Suite	8
3.7	IT Governance and Management	9
3.8	Selected Initiatives	10
4	Economic Dimension	13
4.1	Operational Machine Customers	14
4.2	Predictive Machine Customers	14
4.3	Strategic Machine Customers	14
4.4	Economic Potential and Market Implications	15
5	Political and Regulatory Dimension	16
5.1	Legal Status and Attribution	16
5.2	Civil Liability and Accountability	17
5.3	Data Governance and Data Sovereignty	18
5.4	Financial-Market Regulation and Digital Identity	18
5.5	Implications for Financial Institutions	19
6	Blockchain-based Prototype	20
6.1	Buyer-side Prototype	20
6.2	Provider-side Prototype	22
6.3	Prototype Interpretation	24
7	Conclusion and Outlook	25
	Authors	26
	References	27

1. Introduction

The financial industry is entering a new stage of digital transformation in which non-human actors increasingly participate in economic exchange. Alongside automated processes and digital channels, machines are beginning to act not only as tools that support decision-making, but also as autonomous market participants. In particular, advances in artificial intelligence (AI), connected devices, and digital infrastructures enable autonomous software agents, Internet-of-Things (IoT) devices, and digital twins to make decisions and execute transactions with limited or no human intervention. In this context, Gartner (2025) refers to such actors as “machine customers” and describe them as non-human economic actors that purchase goods or services.

While this establishes the economic role of machine customers, it says less about the capabilities that distinguish them from conventional automated systems. The literature on intelligent agents helps sharpen this distinction. A used notion of agency characterizes such systems by four core properties: autonomy, social ability, reactivity, and pro-activeness (Wooldridge & Jennings, 1995). Applied to machine customers, these properties can be interpreted as follows:

- **Autonomy:** Machine customers can operate without continuous human intervention and exercise some control over their internal state and actions.
- **Social ability:** Machine customers are able to interact with other systems, platforms, service providers, or users through digital interfaces and communication protocols.
- **Reactivity:** Machine customers perceive and process relevant information from their environment and respond to changes in a timely manner.
- **Pro-activeness:** Machine customers do not merely react to inputs, but can pursue goals by initiating actions or transactions.

Understood in this way, machine customers are not simply rule-based automation tools, but systems that can perceive relevant information, make context-dependent decisions, and initiate actions with limited human intervention (Wooldridge & Jennings, 1995; Piccialli et al., 2025). In financial contexts, this means that they may not only ex-

ecute transactions, but also access, consume, and act on financial data on behalf of users, for example by analyzing account information, comparing products, or preparing financial decisions.

For financial service providers, the significance of machine customers extends beyond efficiency gains and new modes of interacting with clients. They may also form a distinct customer segment. Because machine customers can independently initiate and process transactions, including payments, they may act as customers or intermediaries for a broader range of financial products and services, such as investments, lending, and other data-driven offerings (FinRegLab, 2025; OECD & Financial Stability Board, 2024). They may also access and use customer-related and financial data on behalf of users to compare products, support decisions, or trigger follow-up actions. Their emergence challenges providers to rethink how financial services are designed, delivered, and governed in an environment where interactions increasingly take place between machines and financial systems rather than only between human users and institutions. This raises important governance questions regarding data access, authorization, consent, accountability, security, and compliance (FinRegLab, 2025; Swiss Financial Market Supervisory Authority (FINMA), 2024; OECD, 2026).

Against this background, this report examines machine customers primarily from the perspective of financial service providers. The central question is what capabilities, structures, and conditions financial institutions may need in order to become ready for this emerging customer segment. To address this question, the report is structured according to the STEP method. This framework provides a systematic perspective on the social, technological, economic, and political/regulatory dimensions shaping the rise of machine customers and helps to identify the key developments, opportunities, and challenges associated with their growing presence in financial markets.

We would like to thank our research partners Canton of Zug, Finnova, Finstar, Inventx, SFTI / Swiss FinTech Innovations, SIX, and Zürcher Kantonalbank for their financial and content support, which made this report possible.

2. Social Dimension

The social dimension of machine customers concerns the changing relationship between human actors, organizations, and increasingly autonomous digital systems in economic exchange. From the perspective of financial services, this development is relevant because roles that were traditionally reserved for human customers may increasingly be assumed, at least in part, by technical systems acting on the basis of delegated authority, predefined mandates, and data-driven decision processes (Aldasoro, Gambacorta, Korinek, Shreeti, & Stein, 2025; Basel Committee on Banking Supervision, 2024).

From a socio-technical perspective, machine customers should not be understood as isolated technical artifacts. Rather, they are embedded in broader systems involving human principals, organizational governance structures, software architectures, legal accountability, and market infrastructures (Akbarighatar, Pappas, & Vassilakopoulou, 2023; Baxter & Sommerville, 2011). Therefore, adoption is determined less by technical capabilities and more by structural constraints (Heines, 2026). Their significance therefore lies not only in automation itself, but also in the reorganization of exchange relationships between humans, service providers, and digital systems. As such actors become increasingly present in global markets, they may represent not only a technological development but also a new customer segment for financial service providers. This creates opportunities for efficiency gains, revenue generation, and service innovation, while at the same time increasing demands on secure infrastructures, API-based business models, governance, regulation, and risk management (Rhyner, 2025b; OECD & Financial Stability Board, 2024; Swiss Financial Market Supervisory Authority (FINMA), 2024).

For financial service providers, this shift implies that financial services and processes originally designed for human customers must increasingly be adapted to interactions with non-human actors (Luo, Li, Huang, & Zhan, 2024). Machine customers do not primarily interact through relationship managers, mobile applications, or call centers, but through digital infrastructures, standardized interfaces, and machine-readable rules. Banks may therefore need to evolve their business models toward platform-based and API-driven services (Swiss Financial Innovation Desk (FIND), 2025). The provision of secure, scalable, and

governable APIs becomes a strategic competitive factor in such machine-to-business interactions. Because machine customers act on optimization logics, rules, and data rather than emotions or social preferences, pricing, risk assessment, and product design may also need to be reconsidered from the perspective of machine decision-making. At the same time, machine customers may create scope for new products and business models, such as micro-financing for connected devices, AI-supported treasury automation, or trusted wallet services for digital twins (Rhyner, 2025b).

A machine becomes a capable customer when it can act autonomously without continuous human supervision, evaluate options, make decisions, and initiate actions within a defined mandate. In addition, it requires access to resources, budgets, or permissions that it is authorized to use, as well as the ability to interact with systems, services, or other actors in order to execute transactions. It is the combination of autonomy, decision-making capability, resource control, and interactivity that distinguishes machine customers from conventional automated systems (Wooldridge & Jennings, 1995; Piccialli et al., 2025). An illustrative example is provided by the robo-taxi case shown in the box on page 4.

Within the social dimension, machine customers can be differentiated into three ideal types, as shown in Figure 2.1. This typology synthesizes insights from research on autonomous agents, predictive maintenance and optimization, autonomous market decision-making, and automated negotiation (Piccialli et al., 2025; Zhu, Ran, Zhou, & Wen, 2026; Peters, Ketter, Saar-Tsechansky, & Collins, 2013; Luo et al., 2024; Xu, Mak, Minaricova, & Brintrup, 2024). The three types differ primarily in the scope of their autonomy and decision-making capabilities:

- **Operational machine customers:** These systems autonomously execute routine and well-defined transactions, such as reordering consumables or triggering maintenance services. They are primarily associated with automated procurement, sensor-driven decision-making, and rule-based cyber-physical systems.




Figure 2.1: Ideal types of machine customers by increasing autonomy and decision complexity

- **Predictive machine customers:** These systems anticipate future needs and optimize decisions based on data patterns, for example by forecasting energy demand and purchasing accordingly. They are related to predictive analytics, reinforcement learning, and optimization-based decision-making.
- **Strategic machine customers:** These systems operate at a higher level of autonomy by engaging in negotiation, long-term planning, and complex market interactions. Examples include autonomous fleets that negotiate service contracts or optimize multi-party logistics. They are associated with research on agentic AI, multi-agent systems, and autonomous decision-making in dynamic environments.

These three archetypes differ in their degree of autonomy, complexity of decision-making, and implications for governance, risk, and financial service design. They also show that machine-driven economic agency is not confined to a single domain, but may emerge across mobility, energy, logistics, smart-home ecosystems, healthcare, and other digitally connected environments. However, their boundaries may be fluid, as real-world applications can combine features of multiple archetypes or shift between them depending on context, system design, and operational conditions.

For the financial industry, this typology should not be understood as describing a separate sector-specific category of machine customers. Rather, financial service providers may be required to interact with operational, predictive, and strategic machine customers across different application contexts. In this sense, the financial industry occupies a particular position: it is less the primary origin of

these machine customers than an enabling and intermediary domain that must provide the accounts, payment mechanisms, financing solutions, insurance products, and governance structures through which such actors can operate. The relevance of the typology for financial services therefore lies not in assigning the industry itself to one specific type, but in understanding how different forms of machine customers may generate distinct requirements for financial infrastructures, products, interfaces, and control mechanisms.

Illustrative example: Robotaxi as a machine customer

Scenario: A robotaxi autonomously drives to a charging station when its battery falls below a defined threshold and initiates payment for the charging service itself.

Why this is relevant:

- The robotaxi detects a need based on operational data.
- It evaluates available options and selects a provider.
- It initiates and executes a payment without continuous human intervention.
- It acts within predefined rules, permissions, and resource limits.

Broader implication: Similar patterns may also emerge in connected vehicles, production facilities, drones, smart-home systems, or humanoid robots that manage accounts, purchase insurance, initiate payments, or allocate liquidity autonomously. The example illustrates how machines may become actors rather than merely passive technical tools.

3. Technological Dimension

Lead author: Carla Caspar, Inventx AG

The emergence of machine customers, autonomous software agents, connected devices, and AI-driven systems acting on behalf of individuals or organizations, is expected to reshape the IT architecture of banks. Unlike traditional human customers, machine customers do not interact primarily through graphical user interfaces or branch channels. They operate through APIs, events, embedded workflows, and real-time decision loops, and may compare offers automatically, trigger transactions instantly, and optimize outcomes continuously. As a result, banking architectures built primarily for human interaction will increasingly need to evolve toward platforms that are machine-consumable, policy-driven, and capable of operating at digital speed and scale.

To address these requirements, this chapter applies the *Inventx* platform model for a modern IT architecture (Rhyner, 2025a) as a reference model for machine-customer-ready banking platforms.¹ The platform model is illustrated in Figure 3.1 and structures the technical capabilities into seven mutually reinforcing layers: *Digital Experience Platform*, *Decoupling Platform*, *Business Platform*, *Hybrid-Cloud Platform*, *Data Factory*, *Security Suite*, and *IT Governance and Management*. The following sections outline how each architectural layer may need to evolve to enable machine customers, covering machine-readable services, event-driven integration, modular banking capabilities, trust frameworks, and governance models for AI-driven execution.

3.1. Digital Experience Platform

The *Digital Experience Platform* is the external interaction layer through which machine customers discover, evaluate, and consume banking capabilities. In traditional banking, digital experience is primarily designed for humans using mobile apps, web portals, or branch chan-




Figure 3.1: IT architecture blueprint. Source: Based on Rhyner (2025a)

nels. In a machine-customer economy, however, the primary interface shifts from graphical user interfaces to APIs, schemas, events, and self-service onboarding mechanisms. As a result, the quality of the machine-facing experience becomes as important as user experience has been for human customers. Two capabilities are especially critical in this context: first, the design of APIs as high-quality machine-facing products, and second, the provision of API-first portals that enable autonomous actors to discover, access, and integrate banking services with minimal friction.

API UX: Experience Quality for Autonomous Actors

For autonomous agents, an API is the customer interface. Consequently, banks must move beyond merely exposing endpoints and instead design APIs with a strong focus on API user experience (API UX). This means interfaces that are machine-readable, semantically clear, and optimized for automated decision-making. If APIs are inconsistent or poorly documented, automation costs rise and trust declines. Industry research increasingly highlights developer experience, documentation quality, and consistency as critical success factors for API adoption (Gupta, 2025). For machine customers, APIs should therefore be treated not purely as technical infrastructure, but as digital products.

¹ For a broader discussion of future requirements for IT architectures and infrastructures in finance, see Ankenbrand, Bieri, Ettlin, Fischer, and Rhyner (2025).

Developer and Agent Portals (API-First)

A second core capability is the creation of developer and agent portals based on an API-first operating model. Historically, typical bank integrations required lengthy bilateral projects, manual onboarding, and fragmented documentation. Machine customers require the opposite: instant discoverability, self-service access, test environments, and automated trust establishment. Modern portals should therefore provide:

- unified API catalogs for all machine-consumable banking services;
- self-service credentialing for organizations, applications, and autonomous agents;
- usage analytics and service status transparency;
- policy disclosure, including rate limits, consent models, and operational constraints; and
- machine-readable service descriptions, allowing AI agents to discover capabilities automatically.

For future machine customers, the *Digital Experience Platform* becomes the equivalent of today's mobile banking app. The difference is that the "user" is now an autonomous actor that evaluates interfaces based on speed, structure, reliability, and machine interpretability rather than visual design. Banks that provide superior API UX and frictionless self-service onboarding will become preferred partners in automated financial ecosystems.

3.2. Decoupling Platform

The *Decoupling Platform* is the coordination layer between customer-facing channels and core banking capabilities. Its purpose is to separate producers and consumers of services, enabling machine customers to interact with banks in a scalable, secure, and flexible way. To fulfill this role, the *Decoupling Platform* must provide four key capabilities: orchestration across APIs and events, context-aware consent and delegation, standardized agent-to-agent communication, and integration with edge and IoT environments.

Event and API Orchestration

Machine customers trigger actions continuously: requesting credit, rebalancing liquidity, initiating payments, or

reacting to threshold breaches. These interactions cannot depend solely on synchronous API calls. Banks therefore need orchestration capabilities that combine APIs with event-driven architectures. Events such as limit exceeded, invoice received, or foreign exchange rate threshold reached can trigger automated workflows across multiple services in real time. Event-driven architectures reduce dependencies between systems, improve resilience, and support independent scaling of services. This is particularly relevant when thousands of autonomous agents act simultaneously (IBM, online). Regarding API management, error handling mechanisms must be adapted for machine customers. Unlike human users, autonomous agents react deterministically to HTTP status codes and may automatically abort, retry, or escalate transactions without manual intervention. API versioning also becomes particularly important, since machine customers that are configured or trained for a specific API version may not adapt to breaking changes without explicit reconfiguration.

Consent and Delegation Layer

Machine customers should not receive unrestricted access rights. Every action should be executed under a context-specific mandate that defines:

- which machine or agent acts (identity);
- what it is authorized to do (mandate); and
- under which conditions, limits, or time window it may act (context).

This would, for example, ensure that a treasury agent may optimize liquidity but not open a new credit line without additional approval.

A2A Protocols and Standards

As autonomous systems become more common, banks will need standardized agent-to-agent (A2A) communication models. These protocols enable machines to discover capabilities, negotiate tasks, exchange structured requests, and report outcomes securely. Emerging market initiatives such as Google's "Agent2Agent (A2A)" protocol² (Surapaneni, Jha, Vakoc, & Segal, 2025) illustrate how interoperable agent ecosystems may evolve, using capability discovery, task lifecycle states, and asynchronous

² Further selected initiatives are discussed in Section 3.8.

messaging. Such standards could become as important for machine customers as REST APIs are for digital banking.

Edge & IoT Integration

Future machine customers may not operate from central data centers, but from vehicles, industrial systems, smart devices, or local controllers. This requires seamless integration of edge and IoT environments into banking workflows. Edge architectures allow local preprocessing of events, lower latency, and faster autonomous decisions, while cloud platforms provide analytics, governance, and persistent processing. Combining edge and cloud orchestration is increasingly seen as a key pattern for real-time IoT ecosystems (Yende et al., 2025).

3.3. Business Platform

The *Business Platform* contains the core banking capabilities that machine customers consume and combine. In a traditional banking model, services are often bundled into fixed product packages designed for human sales and relationship management. Machine customers, however, require modular, machine-readable components that can be compared, selected, and orchestrated dynamically. This shifts the focus from product monoliths to composable financial services (Scheibenreif & Raskino, 2025). In this context, the *Business Platform* must provide capabilities that can be modularly consumed, priced dynamically, and executed in real time by autonomous actors.

Modular and Machine-Readable Service Components

Autonomous systems do not evaluate branded product bundles. Instead, they assess individual capabilities based on price, performance, constraints, and suitability for a specific task. A treasury agent, for example, may consume separate services such as:

- cash positioning and short-term funding (liquidity service);
- yield optimization on balances (interest service);
- currency conversion and hedging execution (foreign exchange service); and
- payment clearing and transaction finalization (settlement service).

Automated Pricing, Offers, and Timing

Machine customers expect pricing and offers to be generated instantly and continuously. They compare options in real time and select the most advantageous combination. This requires dynamic pricing engines that can calculate rates, fees, limits, and conditions based on market data, customer context, risk profile, and timing. For example, a liquidity agent may choose overnight funding only when rates fall below a threshold, or shift balances automatically when yield conditions improve.

Real-Time Processing

Machine customers operate continuously, not only during business hours. As a result, the *Business Platform* must support real-time execution of decisions, bookings, confirmations, and state changes. Batch-oriented processing models are insufficient when autonomous systems expect immediate responses to events such as payment arrivals, collateral shortfalls, or market movements. Real-time processing enables instant fulfillment, faster decision loops, and improved automation outcomes.

Consequently, the *Business Platform* transforms banking products into intelligent service building blocks. Institutions that provide modular capabilities, dynamic pricing, and real-time execution will be better positioned to serve machine customers that optimize financial decisions continuously and autonomously.

3.4. Hybrid-Cloud Platform

A *Hybrid-Cloud Platform* provides the foundational infrastructure for scalable and resilient execution of machine-driven banking workloads. Unlike human-centric usage patterns, machine customers can generate highly variable and sudden demand spikes, for example when multiple agents simultaneously react to market events or trigger large-scale portfolio rebalancing. This requires infrastructure that can dynamically scale compute, storage, and network resources within seconds. Elastic scalability ensures that banking services remain responsive even under extreme, short-lived workloads. Hybrid-cloud architectures combine private banking environments with public cloud capabilities, enabling institutions to balance regulatory control with virtually unlimited on-demand capacity. This model is widely recognized as a key enabler for modern, distributed financial systems. In addition, intelligent workload management and auto-scaling mechanisms allow systems to allocate resources based on real-time de-

mand signals, which is essential when autonomous agents operate continuously.

3.5. Data Factory

The *Data Factory* is the intelligence layer of the architecture, responsible for producing, structuring, and delivering data in a way that is directly consumable by machine customers. In contrast to traditional data platforms that primarily serve reporting and analytics use cases, the *Data Factory* must support operational, real-time, and decision-driven data consumption. Machine customers require not only access to data, but also precise semantics, explicit constraints, and machine-readable business meaning. To fulfill this role, the *Data Factory* must provide explicit data models, integrate real-time and historical information, and ensure full traceability of machine-driven decisions.

Explicit and Machine-Readable Data Models

A key requirement for serving machine customers is the transformation of product data into explicit, machine-readable models, including product conditions and properties, terms and durations, and additional conditions. Beyond syntactic standardization, semantic standardization becomes equally important to ensure that autonomous systems interpret financial products, events, and actions consistently across different platforms and institutions. In addition, event models become first-class data objects. Events such as limit breaches, liquidity shortages, or foreign exchange threshold triggers are not passive logs, but active signals that can directly initiate machine actions.

Streaming and Historical Data

Machine customers operate on both real-time and historical data simultaneously. For example, a treasury agent may continuously monitor:

- live account balances and transaction streams;
- historical cash flow patterns; and
- market price movements and foreign exchange rate trends.

This requires a data architecture that supports both streaming and batch processing.

Decision Logging and Auditability

A critical requirement in machine-driven banking is full de-

cision traceability. Every action taken by an autonomous agent must be explainable and auditable. This ensures compliance with regulatory expectations and internal governance requirements, particularly in contexts where AI or algorithmic decision-making is involved. Importantly, this approach avoids “black-box automation” by ensuring that decisions are verifiable across time.

3.6. Security Suite

The *Security Suite* establishes the trust and control layer for machine customers operating within banking ecosystems. As autonomous agents increasingly initiate transactions, negotiate services, and execute financial decisions, traditional human-centric security models such as logins, sessions, and manual authentication become insufficient. Instead, security must be continuously enforced at machine speed, with identity, mandate, and context forming the core of every interaction. To support this, the *Security Suite* must establish trusted machine identities, enable secure non-interactive authentication, and provide technical control mechanisms for exceptional or high-risk situations.

Machine Identity and Trust Layer

A foundational capability is machine identity management, often referred to as “Know Your Machine (KYM)”. Every autonomous agent must have a unique and verifiable identity, comparable to how natural persons and legal entities are identified. This identity is enriched with:

- ownership information (who operates the machine or agent);
- mandate definition (what the machine is allowed to do); and
- trust and reputation history (past behavior, reliability, risk signals).

KYM creates a persistent trust layer where machines are not anonymous API consumers but accountable economic actors (Scheibenreif & Raskino, 2025).

M2M Authentication: Non-Interactive and Context-Aware

Authentication in machine-to-machine (M2M) environments must be fundamentally different from human authentication flows. It is:

- non-interactive (no login screens or manual steps);

- sessionless (no long-lived user sessions); and
- cryptographically strong (based on certificates and key material).

Instead of static credentials, machines rely on short-lived, rotating keys and certificates bound to specific workloads or identities. This approach significantly reduces exposure risk and limits the impact of credential compromise. A key enhancement is context-based authorization, where access decisions are not made solely on token validity but on a combination of the machine identity, mandate, and context.

Technical Emergency Mechanisms for Autonomous Agents

Given the autonomous nature of machine customers, the security architecture must also include fail-safe and emergency control mechanisms. These allow banks or governing systems to:

- pause or suspend agent execution in case of anomalies;
- revoke mandates or reduce permissions dynamically;
- redirect or sandbox agent behavior during incidents; and
- enforce circuit breakers for systemic risk scenarios.

Such mechanisms are essential to ensure operational resilience when machine agents operate at high frequency and scale. They act as a “digital kill switch” or containment layer, ensuring that autonomy remains bounded by controllable safety constraints.

3.7. IT Governance and Management

The *IT Governance and Management* layer defines how autonomy, control, and compliance are balanced in a banking environment that increasingly relies on AI-driven and agent-based execution. As machine customers and internal AI agents become active participants in financial decision-making, governance must evolve from static control frameworks into dynamic, machine-readable policy systems that operate in real time. To fulfill this role, the governance layer must define how agents are controlled,

how their actions are monitored, and how policies, regulatory requirements, and compliance obligations³ are enforced in real time.

AI and Agent Governance

A central aspect is the internal governance of AI and autonomous agents. Not every agent is allowed to act in the same way. A clear distinction must, for example, be made between advisory agents, which recommend actions but do not execute them, and execution agents, which are authorized to perform transactions or binding decisions. In addition, the behavior of agents must be versioned and managed like software artifacts, meaning every decision logic, model update, or policy change is traceable over time. This enables rollback, comparison, and controlled deployment of agent capabilities, similar to modern MLOps (Machine Learning Operations) and model lifecycle management practices.

Monitoring and Auditability

In a machine customer ecosystem, observability becomes a core governance requirement. Every action must be fully traceable, including:

- which agent executed the action;
- when the action was initiated, approved, and completed;
- under which mandate and permissions the agent acted;
- which data sources were used; and
- which policy rules were applied at the time of execution.

This creates a continuous audit trail that supports both operational transparency and regulatory compliance.

Policy and Compliance Control

Traditional compliance models rely on manual interpretation and post-hoc validation. In a machine-driven architecture, compliance must be code-based, machine-readable, and continuously enforced. This includes translating regulatory requirements, internal risk policies, and operational constraints into executable rules that can be

³ See Chapter 5 for more detailed information on political and regulatory aspects of machine customers.

evaluated in real time. This approach aligns with emerging concepts such as policy-as-code, where compliance logic is embedded directly into systems rather than applied externally. It ensures that every machine action is automatically checked against applicable constraints before execution, reducing operational risk and increasing consistency in regulatory adherence.

Together, the seven architectural layers describe a machine-customer-ready banking architecture in which external interfaces, integration mechanisms, core banking services, infrastructure, data, security, and governance are designed as mutually reinforcing capabilities. Figure 3.2 summarizes this possible target state by extending the original platform model with the key machine-customer capabilities discussed in the preceding sections.

3.8. Selected Initiatives

The architectural requirements outlined in the previous sections are increasingly reflected in emerging industry initiatives that seek to operationalize machine-customer capabilities. These initiatives do not yet form a fully mature or standardized ecosystem, but they indicate how the technical foundations for machine customers are beginning to move from conceptual architecture toward practical implementation. They address different layers of the emerging machine-commerce stack, ranging from agent-to-service interaction and capability discovery to payment

authorization, agent authentication, and programmable settlement.

One foundational development is the “Model Context Protocol (MCP)”, which aims to provide a standardized interface through which AI agents can discover external capabilities, interpret service interfaces, and interact with software systems under controlled authorization conditions. MCP is not itself a payment protocol. Rather, it addresses a more general infrastructural problem: how autonomous or semi-autonomous agents can reliably connect to external tools, data sources, and services without requiring bespoke integrations for each system. In this sense, MCP can be understood as part of the emerging agent-infrastructure layer that sits beneath machine commerce. It provides a structured mechanism through which agents can identify available functions, understand the context in which those functions may be used, and execute actions subject to predefined access rights and governance constraints (Swiss Fintech Innovations & Acrea, 2026).

This architecture is technologically relevant for machine customers because payment execution is only one component of agentic commerce. Before a machine customer can purchase, negotiate, renew, subscribe, or rebalance a service, it must first be able to discover relevant counterparties, access data, interpret system interfaces, and invoke external capabilities in a controlled manner. Protocols such as MCP therefore illustrate that the future infrastructure of machine customers is likely to depend




Figure 3.2: Machine customer architecture – requirements along the IT architecture blueprint

not only on payment rails, wallets, and digital identity, but also on standardized agent-to-service interaction layers. These layers are important for auditability, permission management, and interoperability, especially where agents operate across heterogeneous enterprise systems, cloud services, financial platforms, and regulated data environments (Swiss Fintech Innovations & Acrea, 2026).

Building on this broader agent-infrastructure perspective, several initiatives focus more specifically on the payment and commerce layer. In September 2025, *Google* introduced the “Agent Payments Protocol (AP2)”, an open protocol intended to enable payments executed by AI agents on behalf of users. *Google* describes AP2 as a common language for secure, compliant transactions between agents and merchants and notes that it was developed with support from more than 60 companies across the payments, technology, and crypto-asset sectors. According to *Google*, AP2 is designed to address three key issues: authorization, that is, proof that the user has granted the agent specific purchasing authority; authenticity, meaning that merchants can verify that the agent reflects the user’s intent; and accountability in cases of error or fraud. *Google* further states that AP2 is payment-method agnostic and is intended to work with different payment instruments, including cards, bank transfers, and stablecoins (Parikh & Surapaneni, 2025).

Google also distinguishes between two usage scenarios:

- **Real-time purchases:** The user is present and the agent supports the transaction in an assisted mode.
- **Delegated tasks:** The user is absent and the agent executes payments autonomously within predefined rules.

This distinction is technologically relevant because it illustrates the shift from assisted digital commerce toward delegated machine commerce, in which authorization, trust, and auditability become core infrastructural requirements rather than optional features (Parikh & Surapaneni, 2025).

Visa has developed a comparable initiative through its “Trusted Agent Protocol”, announced in October 2025. *Visa* presents this as an ecosystem-led framework for AI commerce that is intended to enable secure communication between AI agents and merchants during all stages of a transaction. According to *Visa*, the protocol is designed to help merchants distinguish legitimate AI agents

from malicious bots and to create a more interoperable infrastructure for AI-powered commerce. This points to a technological layer that is especially relevant for strategic machine customers, i.e., trust infrastructures that allow counterparties to authenticate agents and attribute actions within distributed digital environments (Visa, 2025).

A further step toward implementation in a regulated banking environment was announced by *Mastercard* and *Banco Santander* in March 2026. The two firms reported the completion of Europe’s first live pilot end-to-end payment executed by an AI agent within a regulated banking framework. According to *Mastercard* and *Banco Santander*, the pilot transaction demonstrated how AI agents can initiate and complete payments within predefined limits and under existing financial-sector controls (Mastercard, 2026).

A different technical approach is represented by the “x402 Protocol”, an open, HTTP-based payment standard that activates the “402 Payment Required” status code for programmable web payments.⁴ It allows servers to request payment directly within the normal HTTP request-response cycle, eliminating the need for accounts, API keys, or traditional billing systems. When a client requests a resource, the server can respond with a 402 status containing payment instructions encoded in standardized headers. The client then performs an on-chain payment, typically using stablecoins, and retries the request with a cryptographic proof of payment. The server verifies the payment through a facilitator service and, once confirmed, returns the requested resource. This design enables pay-per-use models for APIs, digital content, and machine-to-machine interactions. Because it is chain-agnostic and built on CAIP standards, x402 supports multiple blockchains and can route payments flexibly. Its architecture is intentionally minimal, relying on existing web infrastructure rather than introducing new protocols or wallets (Rhyner, 2026; Reppel et al., 2025; x402 Foundation, online).

Taken together, the architectural layers and selected initiatives discussed in this chapter show that machine customers require more than isolated AI functionality or new payment interfaces. They depend

⁴ The blockchain-based prototype in Chapter 6 applies the x402 protocol to illustrate autonomous service consumption, programmable payment execution, and payment-gated API access.

on a broader machine-commerce stack that combines machine-readable APIs, event-driven integration, modular banking services, hybrid-cloud scalability, real-time data infrastructures, trusted machine identities, policy-based governance, wallets, and flexible payment methods. The emerging industry initiatives around agent access, agent payments, trusted agent authentication, and programmable settlement indicate that these requirements are no longer purely conceptual, but are already beginning to be tested in real financial infrastructures. For banks, this development represents a structural shift to-

ward the integration of machine-readable, event-driven, and policy-based platform architectures. Across all architectural layers, a coherent expansion emerges in which banking functions are provided as granular, semantically unambiguous, and API-native services that can be discovered, evaluated, and consumed by autonomous agents in real time. A key technological success factor for future banks will be the ability to offer trusted, scalable, compliant, and machine-capable decision and execution environments that can operate autonomously and securely for this new customer segment.

4. Economic Dimension

The economic dimension of machine customers concerns how non-human actors reshape market structures, business capabilities, value creation, revenue models, and competitive dynamics. In financial services, this dimension is particularly relevant because machine customers do not merely use digital channels more efficiently than humans, but may also change how demand is generated, how products are compared, and how transactions are initiated and executed. From this perspective, the rise of machine customers is closely linked to questions of market access, distribution, pricing, customer relationships, and the broader reconfiguration of value chains in increasingly automated ecosystems (Gartner, 2026; Sundararajan, Jeena, & Ellis, 2025; OECD, 2025; Rhyner, 2026).

At a general level, the economic impact of machine customers can be described through several interconnected mechanisms:

- **Efficiency gains:** Automated purchasing, optimization, and execution can reduce transaction costs, error rates, coordination frictions, and decision latency, while enabling transactions to be carried out at greater speed and scale.
- **New business models and service capabilities:** Firms are likely to develop products and services specifically for autonomous buyers, including API-based offers, machine-readable contracts, and usage-based services. Value chains are likely to shift toward platforms, data markets, API ecosystems, and orchestration layers that connect autonomous demand with service providers. For machine customers, banks' existing capabilities will need to be expanded. Whereas KYC (Know Your Customer) frameworks are designed primarily for humans, machine customers require new forms of algorithmic trust or KYM (Know Your Machine). Transactions can take place in milliseconds. The ability to profitably provide, conclude, and manage legally compliant, automated contracts for small amounts is therefore becoming essential, together with corresponding fraud detection capabilities (Rhyner, 2026).

- **Changing market structures:** Competition is likely to shift from brand visibility and emotional loyalty toward algorithmic optimization, interoperability, and machine-readable quality signals. Machine customers also generate continuous, granular, and operationally relevant data that can improve forecasting, personalization, and pricing.

These mechanisms suggest that machine customers are not simply an additional sales channel, but a structural economic development. Because autonomous systems can operate continuously on the basis of codified preferences, rules, and data, they may increase market transparency while also intensifying competition around speed, price, reliability, and integration quality. In this sense, economic value increasingly depends not only on the product itself, but also on whether products, services, and processes are accessible to machines through programmable interfaces and standardized information structures (OECD, 2025; Nassr & Kim, 2023).

For financial institutions, this implies that the economic consequences of machine customers extend beyond payment processing. Banks may face new opportunities to diversify customer segments, expand embedded and API-based services, and position themselves within platform ecosystems, but they may also face disintermediation risks if third-party agents control customer interaction and transaction routing. *McKinsey* argues that agentic AI may be particularly disruptive for banking business models that have historically benefited from customer inertia, especially in deposits and credit cards, because AI agents can optimize product choice, cash allocation, and payment routing more systematically than most human customers (Sundararajan et al., 2025). The relevant economic issue is therefore not only whether machine customers exist, but who controls the interfaces, data, trust frameworks, and execution pathways through which they participate in financial markets.

Although operational, predictive, and strategic machine customers share this common economic foundation, they differ in the types of value they create, the market roles they assume, and the implications they carry for pricing, competition, and business-model design, as summarized in Table 4.1.

Table 4.1: Possible economic implications of machine customers by type

Type	Primary value	Basis of competition	Implications for financial institutions
Operational	Efficiency, lower transaction costs	Reliability, machine readability, seamless integration	Stable APIs, low-friction execution, transparent pricing
Predictive	Better timing and resource optimization	Data access, forecasting quality, responsiveness	Opportunities in anticipatory financing, timing automation, and data-driven offers
Strategic	Continuous comparison, negotiation, and demand reallocation	Price, interoperability, settlement terms, algorithmic attractiveness	Disintermediation and customer loyalty risk, margin pressure, and need to remain visible in agent-driven markets

4.1. Operational Machine Customers

Operational machine customers primarily generate economic value through automation and efficiency. Their main contribution lies in reducing the cost of routine transactions, minimizing delays, and lowering the administrative burden associated with repetitive purchasing and payment activities. In financial contexts, this may include automated replenishment, recurring payments, low-value procurement, or predefined service orders. The economic logic is therefore largely based on transaction-cost reduction, process efficiency, and operational continuity. Typical use cases include industrial robotics that autonomously order spare parts or maintenance services, connected vehicles that manage charging or servicing, and IoT devices that trigger procurement processes in real time.

From the perspective of providers, operational machine customers reward machine-friendly service design. In economic terms, this means that machine customer journeys are shaped less by persuasion than by triggers, decision rules, execution conditions, and feedback loops. Availability, reliability, and machine-readable service characteristics therefore become commercially relevant features. Firms that offer stable APIs, reliable execution, transparent pricing, and low-friction payment processes may gain an advantage because such attributes are directly relevant to automated decision systems. In this segment, economic competition is likely to focus less on emotional differentiation and more on execution quality, machine readability, and seamless integration into existing workflows.

4.2. Predictive Machine Customers

Predictive machine customers create value not only by automating execution, but by improving the timing and quality of economic decisions. Because they anticipate

future needs on the basis of continuous data collection and forecasting, they can optimize inventory levels, maintenance cycles, liquidity planning, or procurement timing. In financial-service settings, this may involve dynamic trading and treasury decisions, anticipatory financing needs, automated budget allocation, or the timing of payments in response to predicted future conditions. Illustrative use cases include automated predictive maintenance combined with financing or leasing arrangements, as well as connected systems that reallocate resources or trigger purchases before disruptions occur.

Economically, this shifts the emphasis from pure efficiency toward optimization. The value of predictive machine customers lies in reducing uncertainty, preventing costly disruptions, and allocating resources more effectively over time. This also increases the strategic importance of data access and analytics capabilities, because firms with better forecasting models and richer data may be able to shape demand more effectively and capture more value from machine-mediated transactions.

4.3. Strategic Machine Customers

Strategic machine customers have the broadest economic implications because they are not limited to executing or optimizing predefined tasks, but may actively shape purchasing strategies, negotiate across providers, and reconfigure demand in response to changing market conditions. Their decisions may involve multi-step planning, trade-off analysis, and the comparison of multiple counterparties across price, reliability, financing conditions, and contractual terms. In banking, this could eventually include AI agents acting on behalf of households or firms in seeking financing or liquidity, initiating microcredit requests, managing account or custody relationships, or reallocating funds according to predefined objectives and risk constraints.

This could alter market structures more fundamentally. If strategic machine customers become capable of comparing offers continuously and reallocating demand at scale, competitive advantage may shift away from brand loyalty and toward algorithmic competitiveness. Prices, service levels, settlement conditions, and interoperability could become more important than traditional marketing cues. Agentic AI may therefore erode the “inertia dividend” on which parts of retail and SME banking have long benefited, as agents can, for example, continuously optimize deposits, card usage, and payment execution in line with user objectives (Sundararajan et al., 2025). For financial institutions, this raises the possibility that future competition will increasingly be mediated by agents that optimize across entire ecosystems rather than remaining within a single provider relationship. In such an environment, the economic challenge is not only to serve customers efficiently, but to remain visible, attractive, and interoperable in agent-driven markets.

4.4. Economic Potential and Market Implications

Market analyses suggest that machine customers may become a commercially significant customer segment. *Gartner* reports that CEOs expect machine customers to account for up to 20 percent of revenue by 2030, indicating that firms increasingly view them as a distinct and commercially relevant customer segment rather than a marginal technological curiosity (Gartner, 2026). *McKin-*

sey likewise argues that agentic commerce could become a large-scale economic phenomenon, estimating that the U.S. B2C retail market alone could see up to one trillion dollars in orchestrated revenue by 2030, with global projections of three to five trillion dollars (Schumacher, Roberts, & Giebel, 2025).

For banking, the implications may be especially significant because some revenue pools depend on frictions and forms of customer inertia that autonomous agents could reduce. *McKinsey* identifies deposits and consumer-card economics as particularly exposed areas. AI agents could, for example, monitor balances in real time, shift excess funds into higher-yield accounts, optimize card selection at the point of payment, or route transactions through account-to-account infrastructures where this is economically preferable (Sundararajan et al., 2025). In markets with open-banking access and low-cost instant-payment rails, this could intensify margin pressure, increase price transparency, and weaken the stickiness of existing customer relationships (Sundararajan et al., 2025).

Taken together, these developments suggest that the economic potential of machine customers lies not only in new transaction volumes, but also in their capacity to reshape competitive dynamics. For financial institutions, the challenge is therefore not merely to digitize existing products, but to position themselves within markets in which autonomous systems may increasingly influence product selection, payment routing, and the allocation of financial relationships across providers.

5. Political and Regulatory Dimension

Lead authors: Claude Humbel & Bianca Kremer, Zug Institute for Blockchain Research (ZIBR) & University of Lucerne

The political and regulatory dimension of machine customers concerns the public-policy choices, legal frameworks, and institutional structures that determine whether and how non-human actors may participate in economic exchange and financial markets. In financial services, regulation and legal certainty are preconditions for trust, market integrity, and operational resilience (Swiss Financial Market Supervisory Authority (FINMA), 2024; OECD, 2026; OECD & Financial Stability Board, 2024). Machine customers raise questions about digital identity, legal status, liability, data protection, and risks stemming from financial-market access.

Switzerland traditionally follows a principle-based, technology-neutral, and sector-specific regulatory approach. Current law and regulation thus treat machine customers not as a separate legal category but as instruments whose actions are attributed to the principals on whose behalf they operate. On 12 February 2025, the Federal Council announced that Switzerland would prioritize ratifying the Council of Europe Framework Convention on Artificial Intelligence, focus legislative changes on selected sector-specific adjustments, and encourage non-binding measures to promote responsible AI practices, with a draft legislative package expected by the end of 2026 (Federal Council, 2025a; Wildhaber & Barth, 2025). It is expected that the legislative package will build on the principle-based and technology-neutral approach.

At the EU level, comparable questions are addressed within a broader and more rule-based regulatory architecture, including the AI Act (Regulation [EU] 2024/1689; European Parliament and Council of the European Union (2024c)), DORA (Digital Operational Resilience Act, Regulation [EU] 2022/2554; European Parliament and Council of the European Union (2022)), the Data Act (Regulation [EU] 2023/2854; European Parliament and Council of the European Union (2023)), eIDAS 2.0 (Regula-

tion [EU] 2024/1183; European Parliament and Council of the European Union (2024b)), the GDPR (General Data Protection Regulation, Regulation [EU] 2016/679; European Parliament and Council of the European Union (2016)), and the revised Product Liability Directive (Directive [EU] 2024/2853; European Parliament and Council of the European Union (2024a)). For Swiss financial institutions, these EU developments are relevant where services, clients, infrastructures, or technology providers have a cross-border dimension.

Table 5.1 summarizes these implications by distinguishing operational, predictive, and strategic machine customers and mapping each type to key regulatory questions, primary regulatory anchors, and implications for financial institutions.

5.1. Legal Status and Attribution

Machine customers do not have independent legal capacity. Under Swiss contract law, rights and obligations attach to natural or legal persons. Therefore, declarations or transactions initiated by an automated system must be attributed to the person or organization behind it. Under a doctrine for the attribution of machine declarations (*Maschinenerklärungen*), declarations generated by automated systems are, in principle, imputed to the person who deploys or authorizes the system (Lohmann, 2017, 2024). Where an automated system is used to perform a contractual obligation, the auxiliary-persons rule in Article 101 of the Swiss Code of Obligations provides the relevant doctrinal anchor: the principal is liable for damage caused in the performance of obligations by persons or systems engaged to fulfill them. A growing body of Swiss-law commentary applies the auxiliary-persons rule logic by way of analogy to digital systems, despite their lack of legal personality (Lohmann & Preßler, 2021; Yacoubian, 2023).

For operational and many predictive machine customers, attribution is manageable: the system acts within a defined mandate and can be linked to a registered user, contractual counterparty, or regulated entity. Strategic machine customers may pose more difficult questions. Where agents interact across several providers, negotiate dynamically, or operate through delegated structures or Decentralized Autonomous Organizations (DAOs,

Table 5.1: Political and regulatory implications of machine customers by type

Type	Key regulatory questions	Primary regulatory anchors	Implications for financial institutions
Operational	Scope of mandate, authorization, traceability of routine transactions	Contract law; FINMA AI governance; AML/KYC	Clear delegation rules, auditable transaction logs, embedded payment compliance
Predictive	Data governance, automated decision-making, model risk, forecasting transparency	Data protection acts; FINMA model risk; EU AI Act risk-based duties	Machine-readable consent, model documentation, data minimization, contestation paths
Strategic	Multi-party attribution, liability allocation, agent identity and revocation	Civil/product-liability law; Swiss financial services and financial institutions law; revised EU PLD; eIDAS 2.0; FATF recommendations	Principal-agent structuring, KYM infrastructure, audit trails across counterparties

see Reiser and Wächli (2024)), it might become difficult for counterparties to identify who is legally bound by the agent's actions. Lee (2025) describes the developers, deployers, and infrastructure providers behind an autonomous agent as its "shadow principals", i.e., actors who exert persistent and obscured influence over the agent's actions and challenge the assumptions of agency law. Where attribution fails, a liability gap arises, which in turn fuels proposals to recognize autonomous agents as legal persons. Such proposals have gained little traction in either Switzerland or the EU. The European Parliament's 2017 call for an "electronic personhood" for sophisticated autonomous agents, in particular, was widely criticized as premature and conceptually flawed, and was ultimately abandoned (European Parliament, 2017; Open Letter to the European Commission, 2018; Pullen & Brunner, 2024).

For financial institutions, machine-to-business interfaces and contractual documentation must therefore define the agent's mandate, scope of authority, execution limits, and responsible principal. Emerging protocols address this challenge by encoding verifiable authorization, agent authenticity, and accountability directly into transactions (Parikh & Surapaneni, 2025; Visa, 2025; Swiss Fintech Innovations & Acrea, 2026). These protocols, however, do not determine the legal effect of an agent's action, as that ultimately turns on the law of contract and, in particular, the doctrines of representation and attribution.

5.2. Civil Liability and Accountability

Machine customers can produce harmful outcomes, such as mistaken payments, flawed transactions, unauthorized purchases, or decisions based on incorrect data. *Prima facie*, liability might be difficult to allocate because several actors typically stand behind a single agent: the developer of the AI system, the deployer integrating it into a

service, the operator running it, and the user instructing it (FinRegLab, 2025; Soder, Smakman, Dunlop, & Sussman, 2025; Wildhaber, 2024). Soder et al. (2025) argue that, as agent autonomy increases, responsibility should shift further upstream toward developers and providers best positioned to prevent harm, by analogy with the autonomy-graded liability frameworks already used for automated vehicles.

Switzerland currently addresses AI-related liability through existing rules on civil and product liability, product safety, as well as supervisory frameworks rather than a dedicated AI liability regime. Wildhaber (2024) and Wildhaber and Barth (2025) argue for targeted adjustments to the Product Liability Act and Product Safety Act in light of recent regulatory developments in the EU. These include the revised Product Liability Directive, which will cover software (including AI systems), and the AI Act, which introduces risk-based obligations for certain AI systems and actors. The proposed AI Liability Directive, however, was withdrawn in 2025, leaving no harmonized EU civil-liability regime for AI (European Commission, 2025; European Parliament and Council of the European Union, 2024c; European Parliament, 2026).

For financial institutions, liability risk management therefore depends heavily on governance, contractual risk allocation, and auditability. Institutions need records showing who authorized the agent, which permissions applied, which decisions were taken, and how transactions were executed or revoked. Tibbetts and Jones (2026) flag a similar set of risks across agentic deployments: opacity of decision-making, accountability gaps when agents act autonomously, and the difficulty of evidencing what an agent did and why. Accountability for machine customers is thus as much an organizational question as a legal one.

5.3. Data Governance and Data Sovereignty

Machine customers continuously access, generate, and process data, covering the principal's personal details, transactions, behavior, and potentially third parties' data. In financial services, such data are particularly sensitive: they feed directly into payment authorization, credit assessment, fraud detection, and the design of tailored products.

Swiss and EU law both regulate automated individual decision-making, although with different intensity. Article 21 of the revised Swiss Federal Act on Data Protection requires the controller to inform the data subject of decisions with legal or similarly significant effects taken on a purely automated basis, and to grant, on request, the right to express their views. Article 22 GDPR goes further, establishing a general right not to be subject to such decisions, subject to specified exceptions (Regulation [EU] 2016/679; European Parliament and Council of the European Union (2016)).

Agent-mediated interactions strain these requirements. Consent mechanisms designed for human users do not translate easily into machine-to-machine interactions. Predictive or strategic machine customers may rely on broad and continuous data flows that make data minimization and purpose limitation harder to operationalize. Cross-border operation adds a further layer of complexity: machine customers rely on cloud infrastructures, APIs, payment systems, and data-processing environments distributed across several jurisdictions, raising questions about controller and processor roles, international transfers, and data sovereignty. Beyond data protection, the EU Data Act introduces a complementary regime governing access to and use of operational data generated by connected products and related services, granting users a right to access such data and to share it with third parties of their choice (Articles 4 et seq. Regulation [EU] 2023/2854; European Parliament and Council of the European Union (2023)). For machine customers in financial services, such data are doubly significant: they are economically valuable in their own right and operationally necessary for credit assessment, fraud detection, and the provision of tailored services.

5.4. Financial-Market Regulation and Digital Identity

Financial-market regulation presupposes that customers and counterparties are identifiable and that regulated institutions can comply with prudential, conduct, and anti-money-laundering (AML) obligations. Machine customers complicate this assumption because the technical actor initiating a transaction is not necessarily identical to the legally responsible customer (on the issue in general, see Basel Committee on Banking Supervision (2024) and Swiss Financial Market Supervisory Authority (FINMA) (2024)).

From an AML perspective, institutions must distinguish between the machine customer as technical agent and the natural or legal person behind it. While "Know Your Customer" (KYC) requires the financial intermediary to identify the responsible customer, "Know Your Machine" (KYM) targets the agent acting for that customer and mandates the verification of its identity, scope of authority, technical integrity, and behavioral consistency (Rhyner, 2026). Swiss AML law requires the financial intermediary to identify the contracting party at the beginning of a business relationship and to establish the beneficial owner. Where transactions are routed through automated agents, FATF Recommendation 16 remains relevant: financial institutions must include accurate originator and beneficiary information on wire transfers and ensure that this information remains with the transfer throughout the payment chain (Financial Action Task Force, 2025).

From a prudential and supervisory perspective, the relevant question is not whether the machine customer itself is solvent, since solvency and capital requirements attach to regulated legal persons. The more important issue is whether institutions adequately manage the risks created by machine-mediated activity, including operational, model, outsourcing, cyber, and concentration risk linked to reliance on a small number of agent providers or infrastructures (Aldasoro et al., 2025; Basel Committee on Banking Supervision, 2024). FINMA accordingly expects supervised institutions using AI to maintain appropriate governance, risk management, documentation, and control structures, including identification, assessment, management, and monitoring of AI-related risks, with measures proportionate to the materiality of the applications used (Swiss Financial Market Supervisory Authority (FINMA), 2024).

Against this background, digital identity becomes a central regulatory building block. Machine customers require verifiable agent identity, traceable delegation of authority, and reliable revocation. At the regulatory level, eIDAS 2.0 and the European Digital Identity Wallet provide identity frameworks for natural and legal persons across the EU (Regulation [EU] 2024/1183; European Parliament and Council of the European Union (2024b)). Switzerland's E-ID Act (Federal Council, 2025b) follows a comparable path but currently covers only natural persons, leaving legal persons and non-human actors outside its scope. Werbach (2026) addresses this type of issue, proposing a registration regime for autonomous agents as a new category of legal-governance object, calibrated to risk and combined with smart-contract-based bonding. This proposal illustrates the kind of identity infrastructure that may be needed to extend comparable frameworks to non-human actors.

5.5. Implications for Financial Institutions

For financial institutions, political and regulatory readiness for machine customers is not primarily about recognizing machines as new legal persons. It is about ensuring that existing legal and regulatory functions continue to work when economic activity is initiated or executed by autonomous, non-human systems. Five priorities follow from the analysis above:

- defining who is the responsible principal bound by an agent's actions and within which mandate, scope of authority, and execution limits;
- integrating agent-related risks into compliance, risk management, audit, and operational-resilience frameworks;
- developing KYM capabilities alongside existing KYC and AML processes;
- ensuring that authorizations, disclosures, execution limits, and revocation mechanisms are machine-readable and auditable; and
- engaging with emerging machine-to-machine protocols as components of future regulatory and supervisory infrastructure (Parikh & Surapaneni, 2025; Visa, 2025; Swiss Fintech Innovations & Acrea, 2026).

The rise of machine customers does not require a complete overhaul of existing legal and regulatory frameworks. What is required is the need for those frameworks to remain effective in a more automated environment. The key challenge is to preserve reliable chains of attribution, accountability, and identity, as financial interactions increasingly take place between machines.

6. Blockchain-based Prototype

Lead authors: Denis Bieri, HSLU; Carla Caspar, Inventx AG; Jovana Milojević, HSLU

The preceding chapters developed an overview of machine customers and examined the architectural requirements for a financial system capable of serving them. Building on this conceptual foundation, this chapter introduces a blockchain-based prototype¹ as a concrete technical artifact. The prototype does not imply that future machine-customer interactions must be executed on public blockchains. Instead, the blockchain-based design is used as an experimental sandbox environment for testing autonomous machine-to-machine payments, programmable settlement, and cryptographically verifiable execution. In this sense, the prototype provides a practical instantiation of several requirements identified earlier in the report: machine-readable service consumption, non-interactive authentication, real-time execution, programmable payment handling, and auditability.

The chapter distinguishes between two related prototype designs. The first design, shown in Section 6.1, focuses on the buyer side of the transaction. It demonstrates how an autonomous agent can consume a paid LLM service by reacting to an x402 payment challenge and signing the payment payload with its own wallet. The second design, shown in Section 6.2, extends the perspective to the provider side. It examines how a service provider can expose a payment-protected API service, issue an HTTP 402 Payment Required response, verify and settle the payment through a facilitator, and release the protected service only after the payment condition has been satisfied. Together, the two designs illustrate x402 as a mechanism for linking autonomous service consumption with programmable API monetization.

6.1. Buyer-side Prototype

The first prototype design examines the transaction from the perspective of the machine customer. It focuses on an autonomous agent that accesses a paid LLM service

and completes the required payment without direct human intervention. The purpose is to test whether the agent can technically act as a buyer by receiving a payment challenge, authorizing the payment through its own wallet, and continuing the service interaction after payment execution. The conceptual design of the prototype is illustrated in Figure 6.1. The figure abstracts from the agent's internal reasoning logic and focuses on the payment-mediated service interaction.

The numbered flow depicted in Figure 6.1 consists of the following steps:

1. The autonomous agent initiates an inference request to the paid LLM endpoint.
2. The provider returns an HTTP 402 Payment Required response containing the quoted amount and payment parameters.
3. The agent interprets the challenge, selects an accepted payment option, and uses its programmatic wallet to sign the x402 payment payload locally.
4. The agent retries the original request and attaches the signed payment payload, for example through the PAYMENT-SIGNATURE header.
5. The provider submits the payment information to the facilitator for verification, for example via POST /v2/x402/verify.
6. The facilitator returns the verification result and associated payer metadata.
7. The provider then initiates settlement, for example via POST /v2/x402/settle.
8. The facilitator submits the USDC transfer authorization to the Base blockchain and waits for confirmation.
9. After successful verification and settlement, the provider forwards the request to the selected upstream LLM backend.
10. The provider finally returns the requested output together with the corresponding payment response,

¹ See our GitHub repository here.




Figure 6.1: Blockchain-based prototype for a buyer-side autonomous agent using x402 and USDC on Base to purchase LLM inference. Source: Own illustration

for example through the PAYMENT-RESPONSE header.

This design demonstrates the basic logic of an autonomous machine customer: a software agent identifies a service need, interacts with a machine-readable endpoint through standardized interfaces, receives a price and payment condition, authorizes the payment cryptographically, and receives the requested digital service without human intervention. The prototype therefore combines three elements of agentic commerce: autonomous service consumption, embedded payment execution, and verifiable settlement.

To illustrate the practical feasibility of the buyer-side prototype, a first test is conducted with a funded agent wallet on the Base Sepolia test network. The agent is configured with a private key loaded from an environment variable and uses the *BlockRun* testnet client to call a paid LLM endpoint. Before the request, the wallet balance is queried. The agent then submits a simple prompt to the model `openai/gpt-oss-20b`. The x402 payment flow is executed in the background, the provider returns the LLM response, and the wallet balance is queried again after execution. The resulting execution trace is summarized in the following box.

Illustrative example: Buyer-side agent using x402

Scenario: An autonomous AI agent uses its own funded wallet on the Base Sepolia test network to call a paid LLM endpoint via x402.

Model: `openai/gpt-oss-20b`

Wallet balance before execution: 19.996 USDC

Prompt: *"Hello!"*

Response: *"Hello! How can I help you today?"*

Wallet balance after execution: 19.995 USDC

Cost of execution: approximately 0.001 USDC

Interpretation: The balance is queried before and after the LLM request. The decrease in the wallet balance shows that the agent received the LLM response while the x402 payment was executed automatically in the background.

The same architectural setup can be transferred to the Base mainnet, although the use of mainnet funds introduces additional operational and compliance considerations. The *BlockRun* client is initialized with a private key stored securely as an environment variable. This key corresponds to a self-custodial wallet on the Base network, for example a *Coinbase* wallet or *MetaMask* account. This wallet must be funded with ETH on Base to cover gas fees,

and USDC on Base to pay for API usage. It is at this stage that identity verification becomes relevant: funding the wallet through a centralized exchange such as *Coinbase* requires the user to complete a KYC process and disclose personal information. Once sufficiently funded, the wallet can be used to interact with an LLM via on-chain payments. In the context of machine customers, this implies that any autonomous agent must have programmatic access to a funded wallet in order to consume AI services independently.

6.2. Provider-side Prototype

The second prototype design examines a transaction from the perspective of the service provider. While the buyer-side prototype focuses on whether an autonomous agent can act as a paying machine customer, the provider-side prototype focuses on the reciprocal question of whether a third-party provider can make an API-based resource economically accessible to such agents. The purpose is to test whether a provider can expose a protected API endpoint, define a machine-readable x402 payment condition, return an HTTP 402 *Payment Required* challenge, verify and settle a submitted payment payload through the facilitator, and release the protected service only after successful payment handling.

The provider-side prototype therefore extends the first prototype from autonomous service consumption to programmable service provision. In this setting, the service provider becomes the party that defines the economic access conditions of the resource. It specifies which API route is protected, what amount must be paid, which asset and network are accepted, which receiving address is used, and under which validity conditions the payment authorization is accepted. The provider-side logic then checks incoming requests, issues the payment challenge, receives the signed x402 payment payload, coordinates verification and settlement through the facilitator, and decides whether the request may be forwarded to the protected service logic. The conceptual design of the provider-side prototype is illustrated in Figure 6.2. The figure abstracts from the internal business logic of the protected service and focuses on the payment-mediated access-control process.

The numbered flow depicted in Figure 6.2 consists of the following steps:

1. A customer, application, or autonomous machine agent initiates an API request to the service provider's protected endpoint.
2. The service provider's x402 payment layer checks whether the requested route requires payment and whether the incoming request already contains a valid payment payload for this route.
3. If no valid payment is attached, the provider returns an HTTP 402 *Payment Required* response containing the machine-readable payment conditions for accessing the protected resource, such as the accepted payment scheme, network, asset, amount, receiving address, validity period, and resource identifier.
4. The customer or machine agent retries the original request and attaches the signed x402 payment payload to the API request.
5. The provider submits the signed payment payload and the applicable payment requirements to the facilitator for verification and settlement handling. The facilitator verifies the payment against these requirements, while the provider-side payment layer remains responsible for mapping the request to the relevant protected route.
6. After successful verification and settlement handling, the x402 payment layer unlocks the protected API service and forwards the request to the provider's service logic.
7. The service provider returns the protected service response to the customer or machine agent.

This design demonstrates the basic logic of a provider-side machine-commerce service: a provider exposes a machine-consumable API endpoint, attaches a programmable payment condition to the protected resource, verifies and settles the submitted payment through the x402 infrastructure, and releases the service only after the payment condition has been satisfied. The prototype therefore complements the buyer-side design by showing how x402 can be used not only by autonomous agents that purchase digital services, but also by providers that want to monetize API-based resources at the moment of use.




Figure 6.2: Provider-side view of x402 as an additional payment layer in the API value chain. Source: Own illustration

To illustrate the practical feasibility of the provider-side prototype, a local Express server was implemented off-chain and connected to the Base Sepolia test network for payment verification and settlement. The server exposes a dummy paid API endpoint at `/api/v1/dummy-service` and protects it with x402 middleware configured to require \$0.001 USDC per request. The protected endpoint is intentionally simple: its purpose is not to demonstrate a complex business function, but to isolate the payment-gating logic from the service logic. In the prototype, the endpoint is called without a payment payload, upon which the server returns HTTP 402 Payment Required together with the payment conditions. A test client then retries the request with a signed x402 payment payload. After successful verification and settlement through the facilitator, the server returns HTTP 200 OK and releases the protected response. The resulting execution trace is summarized in the illustrative example box.

From the provider perspective, x402 turns an API endpoint into an economically controlled access point. Instead of relying only on external billing arrangements such as subscriptions, invoices, prepaid accounts, API keys, or platform-based metering, the payment condition is embedded directly into the HTTP request-response cycle. This changes the provider's role in the transaction: the provider no longer only authenticates access or records usage for later billing, but defines a runtime payment condition that must be satisfied before the protected resource is released. The provider can communicate the price and payment requirements at the moment of ac-

cess and link service delivery directly to cryptographic payment authorization and settlement handling (Reppel et al., 2025; x402 Foundation, online).

Illustrative example: Provider-side API using x402

Scenario: A local third-party provider server exposes a dummy paid API endpoint and protects it with x402 middleware. The endpoint is only released after successful payment handling.

Protected endpoint: `/api/v1/dummy-service`

Network: Base Sepolia test network
eip155:84532

Price: \$0.001

Unpaid request: The provider returns HTTP 402 Payment Required and communicates the payment conditions for accessing the protected resource.

Paid retry: After a test client retries the request with a signed x402 payment payload, the provider returns HTTP 200 OK and releases the protected dummy service.

Interpretation: The unpaid request shows that the provider does not release the protected API resource without payment. The paid retry shows that the x402 payment layer can verify and settle the submitted payment through the facilitator before the protected service response is returned.

The HTTP 402 `Payment Required` response functions as a machine-readable payment challenge. It specifies the conditions under which access can be granted, such as the payment scheme, network, asset, amount, receiving address, validity period, and protected resource. This is the provider-side counterpart to the buyer-side wallet authorization described in the first prototype: the buyer-side agent must be able to interpret and sign the payment challenge, while the provider-side endpoint must be able to construct, verify, and enforce it. The paid retry then links API access with cryptographic payment authorization. Before the request is forwarded to the protected service, the provider must ensure that the submitted payment payload corresponds to the correct route, price, asset, recipient, network, and validity window. Verification and settlement form the decision point in this process. The facilitator bridges the HTTP-based service layer and the blockchain-based settlement layer. If the payment is valid and settlement has been accepted or confirmed, the service is released. If verification fails, the provider retains control over the resource and the service remains locked.

The provider-side prototype therefore shows how API monetization can be embedded directly into the technical request-response cycle. Together with the buyer-side prototype, it illustrates both sides of a machine-customer transaction: an autonomous agent can authorize and submit a payment for a digital service, and a provider can define and enforce the payment condition under which that service is released. However, the prototype remains deliberately limited. It uses a local Express server, a dummy service endpoint, a fixed price, and the Base Sepolia test network. It therefore validates the payment-gating mechanism, but not yet a production-grade service offering. In a production setting, the provider would still need to define protected routes, manage dynamic or tiered prices, configure receiving addresses, monitor settlement status, reconcile payments with service delivery, and maintain auditable records of

paid interactions. In regulated financial contexts, this may also raise questions around wallet custody, accounting treatment, stablecoin handling, customer identification, sanctions screening, tax documentation, and integration with internal ledgers. Because machine-customer interactions may occur autonomously and at high frequency, x402 payment handling would also need to be combined with API security controls, including replay protection, timeout handling, duplicate-submission handling, failed-settlement handling, route-level authorization, rate limiting, and abuse prevention.

6.3. Prototype Interpretation

The two prototype designs can be interpreted as minimal illustrations of x402-enabled machine-commerce from both sides of the transaction. They do not aim to demonstrate a complex API product or a complete production architecture. Rather, the buyer-side prototype shows how an autonomous agent can consume a paid LLM service by reacting to a payment challenge, signing a payment payload, and receiving the requested response after payment execution. The provider-side prototype shows how a service provider can expose a machine-consumable endpoint, wrap it with an x402 payment layer, process a paid retry, and release a protected response only after payment handling has succeeded.

Together, the two perspectives illustrate a basic payment-gated API interaction: one autonomous buyer, one machine-readable API service, one programmable payment condition, and one verifiable settlement path. The x402 protocol serves as one concrete instantiation of this pattern. Whether alternative payment protocols, permissioned blockchain networks, or off-chain settlement mechanisms could serve the same function, and under what conditions each would be preferable, remains an open question for further research.

7. Conclusion and Outlook

This chapter presents conclusions and hypotheses based on the report. It summarizes potential implications of machine customers for financial service providers across the social, technological, economic, and regulatory dimensions. The statements should be read as working hypotheses for an emerging field in which technologies, business potential, governance requirements, operating models, and regulatory expectations are still taking shape.

If machines buy services, sellers will also have to become machine-ready. Machine customers are not merely an IT topic or an extension of existing digital channels. They represent an emerging customer segment with distinct interaction patterns, expectations, and requirements. Since they do not rely on branches, relationship managers, call centers, or graphical user interfaces, but interact through APIs, events, schemas, wallets, credentials, and machine-readable rules, financial institutions need products, services, and distribution channels that autonomous agents can discover, evaluate, purchase, and consume. The decisive competitive factor will not only be whether a provider offers digital services, but whether machines can understand, assess, and execute them.

Machine customers create opportunities, but also new competitive pressure. Machine customers may become a new source of demand for payments, data services, financing, liquidity management, investment products, and other financial services. This demand will vary with their degree of autonomy, decision complexity, and functional orientation, ranging from routine execution to anticipatory optimization and strategic decision-making. It opens opportunities for API-based business models, pay-per-use services, embedded finance, and machine-readable products. However, because autonomous agents can compare offers continuously, competition is likely to shift toward price, reliability, speed, and machine-readability. Providers should therefore support aspects such as high data quality, semantic clarity, robust APIs, high availability, low latency, and transparent service conditions.

The new customer segment introduces new risks and governance challenges. As financial interactions shift from humans to autonomous agents, risk management must also evolve. Errors, manipulation, misconfigured

mandates, or malicious behavior can scale rapidly and create operational, financial, or market-level distortions. Financial institutions therefore require, for example, real-time risk frameworks, technical emergency mechanisms, clear mandate structures, and continuous monitoring. Key questions remain around autonomy, accountability, transparency, liability, and the attribution of machine-initiated actions to a responsible human or legal principal.

Trust is the central adoption condition. Trust is required on both sides of the interaction: buyer-side agents must trust providers, and seller-side systems must trust the agents that access their services. This trust cannot be reduced to technical authentication alone. It also depends, among other factors, on identity, authorization, reputation, regulatory compliance, auditability, liability allocation, and revocation mechanisms. For financial institutions, Know Your Customer (KYC) must therefore be complemented by Know Your Machine (KYM).

The prototype demonstrates technical feasibility. The blockchain-based prototype shows that basic machine-commerce interaction is possible. On the buyer side, an autonomous agent can access a paid digital service, react to an x402 payment challenge, sign a payment payload, and continue after payment execution. On the provider side, a service provider can expose a payment-protected API and release the service only after the payment condition is satisfied. However, adoption at scale still requires solutions in areas such as KYM integration, mandate management, fraud detection, dispute handling, and the regulatory treatment of machine-initiated transactions.

Existing legal and regulatory frameworks do not need to be rebuilt from scratch, but they must remain effective in an automated environment. The rise of machine customers does not require a complete overhaul of existing legal and regulatory frameworks. Rather, those frameworks must, where necessary, be selectively adapted and continue to ensure reliable attribution, accountability, and identity when financial interactions are increasingly initiated, negotiated, and executed by machines. This requires clear links to responsible human or legal principals, defined mandates and permissions, and effective liability, oversight, and compliance mechanisms.

Authors

This condensed study was prepared in collaboration with the following individuals, who contributed to its content (in alphabetical order):

Authors HSLU

Thomas Ankenbrand
Head Competence Center Investments

Jovana Milojević
Lecturer

Denis Bieri
Lecturer

Authors Zug Institute for Blockchain Research (ZIBR) & University of Lucerne

Claude Humbel
Assistant Professor

Bianca Kremer
Postdoctoral Researcher

Guest Authors in Addition to the Authors from the HSLU and ZIBR

Sven Biellmann
Working Group Lead Common API
SFTI / Swiss FinTech Innovations

Carla Caspar
Head InventxLab
Inventx AG

Dennis Bilang
Strategic Innovation Manager
Inventx AG

Stephanie Wickihalder
President
SFTI / Swiss FinTech Innovations

We would also like to thank the research partners of this study, namely Canton of Zug, Finnova, Finstar, Inventx, SFTI / Swiss Fintech Innovations, SIX, and Zürcher Kantonalbank, for their monetary and content-related support, including discussions and document reviews.

Contact

For more information about this report, please contact us at:

Thomas Ankenbrand
Lucerne University of Applied Sciences and Arts
thomas.ankenbrand@hslu.ch

Disclaimer

The analyses, models, software elements, or scenarios presented are not intended to serve as a basis for any specific investment, organizational, compliance, or other decisions. They do not substitute for individual professional advice or for legal or regulatory review. Any decision made in reliance on these contents requires an independent assessment, taking into account the relevant facts and circumstances of the particular case. This document includes information obtained from sources believed to be reliable, but the Lucerne University of Applied Sciences and Arts does not warrant its completeness or accuracy. This also includes the outputs of AI tools which were situationally used in the preparation of this document.

References

- Akbarighatar, P., Pappas, I., & Vassilakopoulou, P. (2023). *A sociotechnical perspective for responsible AI maturity models: Findings from a mixed-method literature review*. *International Journal of Information Management Data Insights*, 3(2), 100193. doi: 10.1016/j.jjime.2023.100193
- Aldasoro, I., Gambacorta, L., Korinek, A., Shreeti, V., & Stein, M. (2025). *Intelligent financial system: How AI is transforming finance*. *Journal of Financial Stability*, 81, 101472. doi: 10.1016/j.jfs.2025.101472
- Ankenbrand, T., Bieri, D., Ettlin, J., Fischer, T., & Rhyner, U. (2025). *The Future of Technology in Finance*. Retrieved 28/01/2026, from https://hub.hslu.ch/retailbanking/wp-content/uploads/sites/7/2025/11/IFZ_The_Future_of_Technology_in_Finance-2.pdf
- Basel Committee on Banking Supervision. (2024). *Digitalisation of Finance*. Bank for International Settlements. Retrieved 26/03/2026, from <https://www.bis.org/bcbs/publ/d575.pdf>
- Baxter, G., & Sommerville, I. (2011). *Socio-technical systems: From design methods to systems engineering*. *Interacting with Computers*, 23(1), 4-17. doi: 10.1016/j.intcom.2010.07.003
- European Commission. (2025). *Commission Work Programme 2025 — Annexes (COM(2025) 45 final), Annex IV: Withdrawals*. Retrieved 11/05/2026, from https://commission.europa.eu/strategy-and-policy/strategy-documents/commission-work-programme/commission-work-programme-2025_en
- European Parliament. (2017). *Resolution with Recommendations to the Commission on Civil Law Rules on Robotics (2015/2103(INL))*. European Parliament. Retrieved 11/05/2026, from https://www.europarl.europa.eu/doceo/document/TA-8-2017-0051_EN.html
- European Parliament. (2026). *AI Liability Directive*. Retrieved 11/05/2026, from <https://www.europarl.europa.eu/legislative-train/theme-a-europe-fit-for-the-digital-age/file-ai-liability-directive> (Legislative Train Schedule)
- European Parliament and Council of the European Union. (2016). *Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the Protection of Natural Persons with Regard to the Processing of Personal Data and on the Free Movement of Such Data (General Data Protection Regulation)*. Retrieved 11/05/2026, from <https://eur-lex.europa.eu/eli/reg/2016/679/oj> (OJ L 119, 1–88)
- European Parliament and Council of the European Union. (2022). *Regulation (EU) 2022/2554 of the European Parliament and of the Council of 14 December 2022 on Digital Operational Resilience for the Financial Sector (DORA)*. Retrieved 11/05/2026, from <https://eur-lex.europa.eu/eli/reg/2022/2554/oj> (OJ L 333, 1–79)
- European Parliament and Council of the European Union. (2023). *Regulation (EU) 2023/2854 of the European Parliament and of the Council of 13 December 2023 on Harmonised Rules on Fair Access to and Use of Data (Data Act)*. Retrieved 11/05/2026, from <https://eur-lex.europa.eu/eli/reg/2023/2854/oj> (OJ L, 2023/2854)
- European Parliament and Council of the European Union. (2024a). *Directive (EU) 2024/2853 of the European Parliament and of the Council of 23 October 2024 on Liability for Defective Products and Repealing Council Directive 85/374/EEC*. Retrieved 11/05/2026, from <https://eur-lex.europa.eu/eli/dir/2024/2853/oj> (OJ L, 2024/2853)
- European Parliament and Council of the European Union. (2024b). *Regulation (EU) 2024/1183 of the European Parliament and of the Council of 11 April 2024 Amending Regulation (EU) No 910/2014 as Regards Establishing the European Digital Identity Framework (eIDAS 2.0)*. Retrieved 11/05/2026, from <https://eur-lex.europa.eu/eli/reg/2024/1183/oj> (OJ L, 2024/1183)

- European Parliament and Council of the European Union. (2024c). *Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 Laying Down Harmonised Rules on Artificial Intelligence (Artificial Intelligence Act)*. Retrieved 11/05/2026, from <https://eur-lex.europa.eu/eli/reg/2024/1689/oj> (OJ L, 2024/1689)
- Federal Council. (2025a). *AI Regulation: Federal Council to Ratify Council of Europe Convention*. Schweizerische Eidgenossenschaft. Retrieved 11/05/2026, from <https://www.news.admin.ch/en/nsb?id=104110>
- Federal Council. (2025b). *Switzerland's E-ID Act*. Schweizerische Eidgenossenschaft. Retrieved 11/05/2026, from <https://www.admin.ch/en/e-id-act>
- Financial Action Task Force. (2025). *International Standards on Combating Money Laundering and the Financing of Terrorism & Proliferation: The FATF Recommendations*. FATF. Retrieved 11/05/2026, from <https://www.fatf-gafi.org/en/publications/Fatfrecommendations/Fatf-recommendations.html> (As amended)
- FinRegLab. (2025). *The Next Wave Arrives: Agentic AI in Financial Services*. Retrieved 06/04/2026, from https://finreglab.org/wp-content/uploads/2025/09/FinRegLab_09-04-2025_The-Next-Wave-Arrives-Main.pdf
- Gartner. (2025). *Gartner Unveils Top Emerging Technologies to Support Autonomous Business*. Retrieved 18/04/2026, from <https://www.gartner.com/en/newsroom/press-releases/2025-09-10-gartner-unveils-top-emerging-technologies-to-support-autonomous-business>
- Gartner. (2026). *Machine Customers Are a Trillion-Dollar Opportunity*. Retrieved 18/03/2026, from <https://www.gartner.com/en/publications/when-machines-become-customers>
- Gupta, A. (2025). *API-First Banking: Building For Developers, Not Just Customers*. Forbes Technology Council. Retrieved 13/04/2026, from <https://www.forbes.com/councils/forbestechcouncil/2025/10/16/api-first-banking-building-for-developers-not-just-customers/>
- Heines, R. (2026). *Agentic Wallets – Der nächste Baustein zu autonomen AI-Assistenten*. cc-bei.news. Retrieved 17/04/2026, from <https://cc-bei.news/agentic-wallets-der-naechste-baustein-zu-autonomen-ai-assistenten/>
- IBM. (online). *What is event-driven architecture?* Retrieved 13/04/2026, from <https://www.ibm.com/think/topics/event-driven-architecture>
- Lee, C. (2025). *AI Agents' Shadow Principals*. *U.C. Irvine Law Review*, 17, forthcoming. Forthcoming 2027; GWU Legal Studies Research Paper No. 2026-18; GWU Law School Public Law Research Paper No. 2026-18. doi: 10.2139/ssrn.6358398
- Lohmann, M. F. (2017). *Roboter als Wundertüten — eine zivilrechtliche Haftungsanalyse*. *Aktuelle Juristische Praxis (AJP)*(2), 152–162.
- Lohmann, M. F. (2024). *Wissen, Wille und Erklärungen von Maschinen: Zur Maschinenerklärung im Schweizer und deutschen Privatrecht*. Zürich: Schulthess.
- Lohmann, M. F., & Preßler, K. (2021). *Die Rechtsfigur des Erfüllungsgehilfen im digitalen Zeitalter: Ein deutsch-schweizerischer Rechtsvergleich*. *Recht digital (RD)*(11), 538–547.
- Luo, X., Li, Y., Huang, Q., & Zhan, J. (2024). *A survey of automated negotiation: Human factor, learning, and application*. *Computer Science Review*, 54, 100683. doi: 10.1016/j.cosrev.2024.100683
- Mastercard. (2026). *Santander and Mastercard complete Europe's first live end-to-end payment executed by an AI agent*. Retrieved 31/03/2026, from <https://www.mastercard.com/news/europe/en/newsroom/press-releases/en/2026/santander-and-mastercard-complete-europe-s-first-live-end-to-end-payment-executed-by-an-ai-agent/>
- Nassr, I. K., & Kim, H. (2023). *Shifting from open banking to open finance: results from the 2022 OECD survey on data sharing frameworks*. OECD Business and Finance Policy Papers (No. 24). doi: 10.1787/9f881c0c-en

- OECD. (2025). *Artificial intelligence and competitive dynamics in downstream markets*. OECD Roundtables on Competition Policy Papers (No. 331). doi: 10.1787/ccf0624a-en
- OECD. (2026). *Supervision of artificial intelligence in finance: challenges, policies and practices*. OECD Artificial Intelligence Papers (No. 54). doi: 10.1787/92743dc1-en
- OECD, & Financial Stability Board. (2024). *OECD-FSB Roundtable on Artificial Intelligence (AI) in Finance: Summary of Key Findings*. OECD and Financial Stability Board. Retrieved 01/04/2026, from <https://www.fsb.org/uploads/OECD-%E2%80%93FSB-Roundtable-on-Artificial-Intelligence-AI-in-Finance.pdf>
- Open Letter to the European Commission. (2018). *Open Letter to the European Commission: Artificial Intelligence and Robotics*. Retrieved 11/05/2026, from <https://robotics-openletter.eu/>
- Parikh, S., & Surapaneni, R. (2025). *Announcing Agent Payments Protocol (AP2)*. Google Cloud. Retrieved 31/03/2026, from <https://cloud.google.com/blog/products/ai-machine-learning/announcing-agents-to-payments-ap2-protocol>
- Peters, M., Ketter, W., Saar-Tsechansky, M., & Collins, J. (2013). *A reinforcement learning approach to autonomous decision-making in smart electricity markets*. *Machine Learning*, 92, 5–39. doi: 10.1007/s10994-013-5340-0
- Piccialli, F., Chiaro, D., Sarwar, S., Cerciello, D., Qi, P., & Mele, V. (2025). *AgentAI: A comprehensive survey on autonomous agents in distributed AI for industry 4.0*. *Expert Systems with Applications*, 291, 128404. doi: 10.1016/j.eswa.2025.128404
- Pullen, J., & Brunner, S. (2024). *Rise of the Robots: Können und sollten Roboter rechtsfähig sein?* *ex/ante*(2), 115–127.
- Reiser, N., & Wächli, B. (2024). *Rechtliche Erfassung von DAOs – Teil 1*. *Schweizerische Juristen-Zeitung (SJZ)*(18), 755–769.
- Reppel, E., Caspers, R., Leffew, K., Organ, D., Kim, D., & Dalal, N. (2025). *x402: An open standard for internet-native payments*. Coinbase Developer Platform. Retrieved 06/04/2026, from <https://www.x402.org/x402-whitepaper.pdf>
- Rhyner, U. (2025a). *Die Zukunft der IT-Architektur von Banken und Versicherungen - Der Weg zur innovativen, skalierbaren & resilienten Finanz-IT*. Inventx. Retrieved 13/04/2026, from https://www.inventx.ch/wp-content/uploads/2025/12/Inventx-Booklet_Composable-Business_new.pdf
- Rhyner, U. (2025b). *Machine Customers: Wenn Maschinen zu Kunden werden*. Inventx. Retrieved 16/03/2026, from <https://www.inventx.ch/blog/machine-customers-wenn-maschinen-zu-kunden-werden/>
- Rhyner, U. (2026). *Machine Customers: Die Business-Architektur einer Bank*. Inventx. Retrieved 06/04/2026, from <https://www.inventx.ch/blog/machine-customers-die-business-architektur-einer-bank/>
- Scheibenreif, D., & Raskino, M. (2025). *When Machines Become Customers: Ready or not, AI enabled non-human customers are coming to your business. How you adapt will make or break your future*. Gartner. Retrieved 13/04/2026, from <https://www.gartner.com/en/publications/when-machines-become-customers>
- Schumacher, K., Roberts, R., & Giebel, K. (2025). *The agentic commerce opportunity: How AI agents are ushering in a new era for consumers and merchants*. McKinsey & Company. Retrieved 31/03/2026, from <https://www.mckinsey.com/capabilities/quantumblack/our-insights/the-agentic-commerce-opportunity-how-ai-agents-are-ushering-in-a-new-era-for-consumers-and-merchants>
- Soder, B., Smakman, J., Dunlop, C., & Sussman, O. (2025). *An Autonomy-Based Classification: AI Agents, Liability and Lessons from the Automated Vehicles Act*. Retrieved 11/05/2026, from <https://www.interface-eu.org/publications/ai-agent-classification>

- Sundararajan, R., Jeenah, U., & Ellis, A. (2025). *The end of inertia: Agentic AI's disruption of retail and SME banking*. McKinsey & Company. Retrieved 31/03/2026, from <https://www.mckinsey.com/industries/financial-services/our-insights/the-end-of-inertia-agentic-ais-disruption-of-retail-and-sme-banking>
- Surapaneni, R., Jha, M., Vakoc, M., & Segal, T. (2025). *Announcing the Agent2Agent Protocol (A2A)*. Google Developers Blog. Retrieved 13/04/2026, from <https://developers.googleblog.com/en/a2a-a-new-era-of-agent-interoperability/>
- Swiss Financial Innovation Desk (FIND). (2025). *Building Blocks for Smart Finance in Switzerland*. Retrieved 25/03/2026, from <https://www.sif.admin.ch/dam/de/sd-web/0uGU5ZH8-FDO/Buliding%20Blocks%20for%20Smart%20Finance%20in%20Switzerland.pdf>
- Swiss Financial Market Supervisory Authority (FINMA). (2024). *Governance and Risk Management when Using Artificial Intelligence* [Guidance 08/2024]. Retrieved 01/04/2026, from <https://www.finma.ch/en/news/2024/12/20241218-mm-finma-am-08-24/>
- Swiss Fintech Innovations, & Acrea. (2026). *Reference Architecture for AI Agent Access to SFTI APIs* [SFTI project]. Retrieved 06/04/2026, from <https://swissfintechinnovations.ch/wp-content/uploads/2026/02/SFTI-White-Paper-Reference-architecture-for-AI-agent-access-to-SFTI-APIs.pdf>
- Tibbetts, S., & Jones, S. (2026). *The Agentic AI Revolution: Managing Legal Risks*. Squire Patton Boggs. Retrieved 11/05/2026, from <https://www.squirepattonboggs.com/insights/publications/the-agentic-ai-revolution-managing-legal-risks/>
- Visa. (2025). *Visa Introduces Trusted Agent Protocol: An Ecosystem-Led Framework for AI Commerce*. Retrieved 31/03/2026, from <https://usa.visa.com/about-visa/newsroom/press-releases.releaseId.21716.html>
- Werbach, K. (2026). *Agents, Inc.* SSRN working paper. doi: 10.2139/ssrn.6465238
- Wildhaber, I. (2024). *KI und Haftung: Lösungsansätze für die Schweiz*. Jusletter IT.
- Wildhaber, I., & Barth, F. (2025). *Regulierung der künstlichen Intelligenz in der Schweiz*. Aktuelle Juristische Praxis (AJP)(6), 598–614.
- Wooldridge, M., & Jennings, N. R. (1995). *Intelligent agents: theory and practice*. *The Knowledge Engineering Review*, 10(2), 115–152. doi: 10.1017/S0269888900008122
- x402 Foundation. (online). *Core Concepts - HTTP 402*. Retrieved 06/04/2026, from <https://docs.x402.org/core-concepts/http-402>
- Xu, L., Mak, S., Minaricova, M., & Brintrup, A. (2024). *On implementing autonomous supply chains: A multi-agent system approach*. *Computers in Industry*, 161, 104120. doi: 10.1016/j.compind.2024.104120
- Yacoubian, C. (2023). *Digitale Systeme als "Erfüllungsgehilfen": Relevanz der fehlenden Rechtsfähigkeit?* Aktuelle Juristische Praxis (AJP)(4), 412–422.
- Yende, R. G., Mambani, G. M. W., Kalombo, M. M.-S., Parfum, B. C., Tabiaki, T. R., & Mumbere, S. M. (2025). *Architectural Integration of Event-Driven Paradigms within Intelligent Internet of Things Ecosystems*. *Scientia. Technology, Science and Society*, 2(12), 32–58. doi: 10.59324/stss.2025.2(12).03
- Zhu, T., Ran, Y., Zhou, X., & Wen, Y. (2026). *A Survey on Intelligent Predictive Maintenance (IPdM) in the Era of Fully Connected Intelligence*. *IEEE Communications Surveys & Tutorials*, 28, 633-671. doi: 10.1109/COMST.2025.3567802

**Lucerne School of
Business**
Institute of Financial
Services Zug IFZ
Campus Zug-Rotkreuz
Suurstoffi 1
6343 Rotkreuz

T +41 41 757 67 67
ifz@hslu.ch
hslu.ch/ifz

A study conducted by

HSLU Lucerne University
of Applied Sciences
and Arts



ISBN-Number
978-3-907379-69-1