GLEIF

Enabling global identity
Protecting digital trust

Agentic AI in Payments: Establishing Interoperable Trust and Control

Working Paper Series

September 2026

Disclaimer

Acknowledgements

This paper is a collective undertaking of multiple contributors including Kevin Griffin, Karla McKenna, Ivan Mortimer-Schutts, and Aydar Negimatzhanov (GLEIF), Petar Atanasovski and Lazar Travica (Curvy Protocol), Miroslav Milenkovic (independent), Mai Fujimoto (zERC20) and Gary Liu and Amit Arora (Terminal 3).

The paper was informed by a workshop convened in April 2026 by GLEIF and Curvy along with participation and notable contributions from experts including Ognjen Kurtic (Finspot), Marija Vlajković (World Bank), Lav Odorovic (Relio), Thomas Hardjono (MIT), Maxim Nescheret (CMA Small Systems AB), and Chris Ostrowski (SODA). The content has also been revised based on inputs received during the Point Zero Forum Roundtable “Solving for the Fundamentals: Identity Wallets, Trust Infrastructure, and Preparing for the Agentic Economy”, held in Zurich on June 23rd and the subsequent expert discussions hosted by the Cardano Foundation and GLEIF.

About GLEIF Working Paper Series

GLEIF publishes working papers and white papers on topics relevant to its mandate and the development of digital organizational identity as a foundation or element of digital public infrastructure. The papers are meant to inform debate and discussion on topics of policy, governance, standards, operational, and technical issues and bring together diverse perspectives from across the global ecosystem.

These papers may be authored by its staff or by and with contribution of third parties. The content of these should not be taken as an indication of the current or future views or policies of the GLEIF Board or its Regulatory Oversight Committee, nor do they necessarily represent official positions of GLEIF management. They are issued as contributions to collegiate discussion of policy, standards and operational issues in identity management, finance and the digital economy.

Executive Summary

Agentic AI is moving from experimental applications to deployment in financial services and digital commerce. Major payment networks, technology providers and financial institutions are introducing infrastructure that enables AI Agents to discover services, negotiate access to data, initiate transactions and execute payments with increasing levels of autonomy. This evolution represents a transition from user-interface-driven commerce towards machine-mediated economic activity, creating new requirements for trust, governance and interoperability.

While Agentic AI does not fundamentally challenge existing legal principles, it raises questions about how it should be controlled. Experts raise technical, operational and commercial questions about how autonomous software acting on behalf of individuals, businesses and public institutions should be authorized, controlled and supervised. Existing legal frameworks are capable of allocating responsibility and liability, but practical mechanisms are needed to demonstrate delegated authority, establish non-repudiation, enforce organizational policies, protect consumers and provide reliable evidence where disputes arise. The challenge is therefore less one of creating new law than of translating established legal principles into robust technical and operational practice.

This paper examines the need for identity and trust infrastructure for Agentic payments. AI Agents increasingly operate across organizational, technological and jurisdictional boundaries, interacting not only with humans but also with other Agents, payment providers and data services. To function safely at scale, these ecosystems require more than digital identity alone. They require mechanisms to establish verifiable organizational identity, express and manage delegated authority, enforce machine-readable policies at runtime, generate auditable records, and support privacy-preserving yet compliant execution across multiple platforms.

The paper explores how the Global Legal Entity Identifier (LEI) ecosystem and the verifiable LEI (vLEI) could contribute to this emerging trust architecture. It posits the need for and discusses the design of Agent Mandate Credentials (AMCs), enabling AI Agents to present cryptographically verifiable proof of identity, delegated authority and organizational relationships. Rather than proposing a single technical architecture, the paper uses practical payment use cases—including treasury automation, payment for data, cross-border settlement and multi-Agent workflows—to identify the functional capabilities that future infrastructures will require.

Scaling control of Agentic AI requires agile, verifiable and highly bespoke management of delegated authority. AI Agents seem likely to work in groups, for many sub-tasks and limited lifespans, and for particular contractual arrangements. Parties responsible for

agents will therefore need to issue, manage, monitor and revoke credentials or their policies on a frequent basis and even potentially "on the fly". Yet their links to other parties will need to be maintained to assess legitimacy and the actors behind them. This may pose challenges for credential ecosystems that rely at each step on new third-party issuance and attestation processes.

Successful governance of Agentic AI will also depend on economic sustainability. Trust frameworks must create incentives for adoption, reduce transaction and compliance costs, support commercially viable business models and avoid unnecessary market fragmentation or platform lock-in. The economics of digital identity, delegated authority and verification therefore deserve the same attention as their legal and technical design if Agentic AI is to scale securely, efficiently and with broad market confidence.

Collaboration among different ecosystems stakeholders is needed to design and test the inter-related technical, economic, operational and legal aspects of frameworks for controlling Agentic AI. Policy makers and standard setters should facilitate practical testing and development by private sector parties around pertinent use cases. This paper outlines some of the practical issues and opportunities raised with GLEIF in multi-stakeholder industry roundtables and workshops. These could be built upon as examples through which to enhance common understanding of operational requirements and solutions and contribute to standards that industry can deploy.

Preface

Recent developments in Agentic payment infrastructure demonstrate that the market is moving rapidly from experimental AI-assisted workflows toward autonomous machine-mediated commercial interactions. Between 2025 and early 2026, major payment networks and infrastructure providers including Mastercard, Visa, Stripe, Google and Coinbase-backed X402 initiatives launched dedicated infrastructure for Agentic payments. These developments indicate that questions of identity, authorization, delegation, auditability and policy enforcement are no longer theoretical.

AI Agents are becoming increasingly important actors in digital services and payments. They extend AI applications beyond generating responses to user prompts, enabling systems to manage complex workflows autonomously. These workflows may result in the execution of economic and contractual obligations, including transactions and payments, carried out on behalf of individuals, companies, and public institutions.

AI Agents introduce a new level of system complexity and autonomy. They do not remove responsibility from the organizations that deploy them. However, they increase reliance on well-defined policies, decision frameworks, and permission systems. These must govern interactions between autonomous Agents that operate according to defined goals. Their use raises legal and compliance questions, particularly in relation to delegation of authority, transaction authorization, non-repudiation and fraud prevention.

The rise of Agentic AI introduces a transition from user-interface-driven commerce to machine-readable, API-native economic interaction. AI Agents increasingly discover services, negotiate access to data, procure digital resources, and initiate payment actions without direct human intervention at each step. This creates new requirements for cryptographically verifiable delegation and policy enforcement.

The use of AI Agents in payments also creates practical challenges related to identity, digital signatures, and access control. Legal responsibility depends on clear identification of the parties that consent to transactions or enter into agreements. Proof of identity and consent therefore remains essential to manage legal, compliance, and operational risks. This issue is highlighted in a January 2026 publication by the Infocomm Media Development Authority (IMDA) of Singapore, *Model AI Governance Framework for Agentic AI*, which states:

> *"While existing governance principles for trusted AI such as transparency, accountability and fairness continue to apply, they need to be translated into practice for Agents."*

This paper contributes to the ongoing dialogue among policymakers, standards bodies, financial institutions, payment infrastructures, AI developers and digital identity providers. It examines how existing legal principles and payment governance frameworks may be operationalized in increasingly autonomous payment ecosystems. Without a robust trust infrastructure, risks include (i) unauthorized transactions, (ii) fraudulent Agents, (iii) compromised systems, and cascading failures across critical infrastructures such as payments, energy, and supply chains. This paper sets out a framework for development of mechanisms that enable cryptographically provable intent and mandates for AI Agents.

The paper also provides examples of use-cases and potential implementation approaches. It outlines how the GLEIF ecosystem and verifiable Legal Entity Identifier (vLEI) could support new Agent Mandate Credentials (AMCs) that enable the functions or rights delegated to AI Agents to be managed in policies and their interaction with third parties and other AI Agents to be controlled.

Table of contents

1 What is Agentic AI? 9
2 Policy and Legal Considerations – a Literature Review 12
3 Policy Concerns and Risks 16
4 Use Case Applications and Requirements 18
5 Operational Challenges: From Principles to Practice 30
6 Design Principles, Options and Open Issues 35
7 Conclusions, Requirements and Approaches 41
Annex A: vLEI-Based Agent Mandate Framework 44
Annex B: PILOT CONCEPT 1: Supplier payments and treasury automation 49
Annex C: PILOT CONCEPT 2: Payments for data for due diligence 52
Annex D: The Economics of Verification Models 54
References 58

Box 1: Integration with EU Business Wallet Architecture 17
Box 2: Requirements for controlling Agentic AI 26
Box 3: Privacy as a functional requirement 28
Box 4: Flexible attribute layer management design options 36
Box 5: Practitioner example of decoupling reasoning and execution 38
Box 6: The GLEIF ecosystem and application to identification and control of AI Agents 48

Figure 1: Agentic AI relationships (illustrative) 9
Figure 2: Composite agentic identity ecosystem actors 34
Figure 3: Architecture for an Agent Mandate Credential 45
Figure 4: Use Case summary for payments automation 49
Figure 5: Use case synopsis for payments for due diligence data 52

1 What is Agentic AI?

This paper draws on related literature and defines Agentic artificial intelligence (AI) as a class of tools that users can assign to perform complex tasks. These systems are given a degree of autonomy to determine how tasks are executed.

At this stage, there is no apparent universally accepted definition of Agentic AI. However, it differs from generative AI in an important way. Generative AI requires users to prompt and act on outputs step by step. In contrast, Agentic AI systems can act independently, interact with other Agents, and execute transactions on behalf of users. In this model, users delegate authority to an AI Agent within defined boundaries.

AI Agents operate within a broader ecosystem rather than in isolation. They depend on supporting infrastructure such as sandboxes, identity systems, and policy engines. These components control access to data and define how Agents operate. Agents act within permissions set by users and governed by enterprise policies. Standards bodies such as ETSI describe AI Agents as autonomous systems that can collect data, learn from past experience, and improve decision-making to perform specific tasks[1]. NVIDIA emphasizes that AI Agents rely on an array of tools and infrastructure including sandboxes, identity controls, and policy engines to secure access to and protect data. These Agents operate in accordance with permissions and goals set by end users and framed by policies of enterprises

Figure 1: Agentic AI relationships (illustrative)

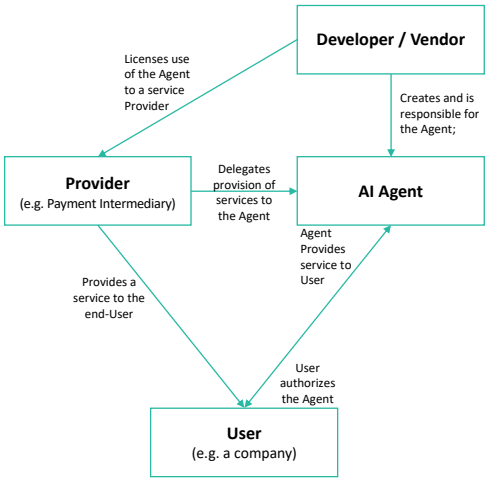

Source: Author's own illustration

AI Agents also exist within a network of legal and contractual relationships. In some cases, an organization deploys an Agent to serve its customers. In others, users select and authorize their own Agents. In all cases, the Agent acts within a framework of delegated authority. The nature of these relationships is central to the issues discussed in this paper. Figure 1 outlines basic relationships in a scenario under which a firm employs an Agent to service the needs of an end-user (in this case a company) to perform specific tasks. The provider delegates services to the Agent in this illustration. But one could also consider a scenario in which the end user or consumer chooses their own Agent and delegates authority to it for the performance of specific tasks. Issues presented in this

[1] ETSI GR ENI 051 V4.1.1 (2025-02)

paper relate extensively to the nature and context of these relationships in which an AI Agent is authorized to act on behalf of a given user.

Example Applications of Agentic AI in payments

Payments and payment services have long been subject to high levels of standardization and automation. The volume of payments processed daily for corporates, consumers and governments requires that transactions are already initiated, verified, authenticated and settled using highly automated processes. Humans remain in the loop however to specify transaction details and check for errors.

Agentic AI can bring this to a further level of automation by executing payments. It can, for instance on behalf of a consumer, automate the full process from discovery and selection to purchase and payment based on pre-set goals and criteria approved by the user. The Agent may transfer secure payment information to authenticate and approve a transaction, for instance with a merchant (for a card payment) or with a bank (e.g. for a credit transfer). There are several examples of firms that have now published protocols and integrated commercial services to apply Agentic AI to different payment contexts. Many e-commerce and digital service providers now offer use of AI Agents to perform complex, delegated transactions. A brief scan of the market highlights some tangible examples.

ChatGPT itself offers an Agentic commerce protocol in partnership with Stripe. Their merchant solution helps simplify consumer experience by delegating the actual payment of a product to their "Instant Checkout", without the need for the consumer to leave the chat. Customers in effect can delegate the payment authorization to ChatGPT, removing the need for the consumer to manually enter and confirm payment details themselves.

Amazon presents a framework to handle more complex and varied payment automation. Mattoo, Nandwani and Villegas outline a sophisticated multi-Agent architecture for their "Cognitive Payments Director" (CPD) Agent. They outline an approach in which distinct Agents are assigned different tasks within an overall framework to address compliance, analyze different payment options and channels, monitor operational status and evaluate performance and commercial terms of different Payment Service Providers (PSP).

Payment card networks Visa and Mastercard have developed solutions to manage Agentic payments. Visa Intelligent Commerce (launched April 2025) and Visa Trusted Agent Protocol (October 2025, with Cloudflare, Microsoft, Shopify, Stripe, Worldpay), have been processing secure Agent-initiated transactions in production since December 2025, as has Mastercard Agent Pay (2025) with "Agentic tokens" and Verifiable Intent.

Google announced the launch of an Agent Payments Protocol (AP2) in 2025, developed with payment companies to enable consumer to merchant Agent-led payments across

platforms[2]. They highlight the need for autonomous Agents in payments to address authorization (proving an Agent had authority), authenticity (to allow a merchant to verify the user's intent) and accountability (to determine who is responsible in the case of fraud or error). It operates by establishing mandates to either present and confirm individual transactions or to execute them based on pre-set rules confirmed by the user. Their intent is to establish a non-repudiable audit trail from intent to payment.

Aldasoro and Desai (BIS 2025)[3] also provide examples of how Agentic AI can be applied to advanced cash management and treasury payments. Building on Gen AI models, they explore how Agentic AI can bring not only benefits through automation but also, by learning over time, can adapt to evolving market conditions to, for instance, optimize liquidity and manage risk. Their findings suggest that significant routine cash management functions can viably be automated using Agentic AI but note that authorities may need to consider policy safeguards to protect against risks in an "era of AI-driven payment operations".

As of 3 June 2026, **public reporting indicates that Stripe, Visa, and Mastercard are collaborating on a joint stablecoin infrastructure platform, with Coinbase considering participation**. Mastercard the same week announced expansion of always-on stablecoin settlement capabilities. Visa's stablecoin settlement pilot expanded to nine blockchains in April 2026. The consolidation of stablecoin infrastructure under major payment networks reduces fragmentation but raises concentration questions.

[2] https://cloud.google.com/blog/products/ai-machine-learning/announcing-agents-to-payments-ap2-protocol
[3] Aldasoro, Iñaki and Desai, Ajit. AI Agents for cash management in payment systems. BIS Working Papers No 1310, November 2025.

2 Policy and Legal Considerations – a Literature Review

Anticipating widespread and more sophisticated use of AI Agents, policymakers, legal and industry experts have begun to assess their impact and consider measures to address implications. Government authorities have reviewed whether AI Agents create fundamentally new risks or whether they can be accommodated within existing legal frameworks and regulation. Prevalent legal concerns relate to attribution of responsibility for Agents that exercise significant autonomy, the adequacy of existing contract law and how to protect consumers without unduly curtailing innovation. Legal experts are anticipating the kinds of claims that may arise and how they would be treated under the law. And technology and financial industry players have also begun to consider practical implications for the adjustment of business and operations, including technical mechanisms to control and audit AI Agents. This section summarizes key arguments from recent literature.

The IMDA (2026)[4] notes in its recent publication on Agentic AI that "While existing governance principles for trusted AI such as transparency, accountability and fairness continue to apply, they need to be translated in practice for Agents. Meaningful human control and oversight need to be integrated into the Agentic AI lifecycle." This position highlights the focus of this discussion document: what are the practical issues and mechanisms we need to consider, if we are to put existing principles into practice for AI Agents? Do we have the right tools and how can they be inserted more effectively into a complex, quickly evolving and international setting for Agentic AI?

The Australian Productivity Commission (APC) report on "Harnessing data and digital technology" solicited feedback from industry on reforms needed to enhance public confidence in the role of AI. In the context of discussion about the Consumer Data Right (CDR), they highlighted the impact that AI Agents are likely to play in accelerating (demand for) access to data, both standardized and unstructured, from across platforms and product providers to inform payments. This will further boost demand and market development of data intermediaries and third parties and implicitly strengthen the importance of solutions that empower consumers with greater transparency and control[5].

The APC also discusses the implications of Agentic AI for the development of data intermediaries and third parties. They note that these types of service providers are developing solutions that allow consumers "to bring together non-standardized or unstructured data that is available from the products and platforms they use. This has the

[4] Infocomm Media Development Authority (IMDA) of Singapore. MODEL AI GOVERNANCE FRAMEWORK FOR AGENTIC AI Version 1.0 | Published 22 January 2026

[5] See the Report: Harnessing data and digital technology | Inquiry report No. 111 | 10 December 2025, Australian Productivity Commission.

potential to support more effective data access and use over time. But it also increases the importance of being able to control and track provenance of data, especially where models still need to be 'explainable' and to enable robust auditability in the case of litigation.

The Digital Regulation Cooperation Forum (DRCF) of the United Kingdom highlighted the need to identify who controls Agentic AI, who is accountable and the challenges for consumers. Their consultation and "foresight report" collated inputs from private and public stakeholders on the risks and opportunities associated with agentic AI. They note that the advancement of AI "from tool to actor" requires increased coordination, which in turn also requires effective identity and access management. These are commonly provided for via technical controls set up by deployers themselves but this can create vendor lock-in. Governance considerations need to be a focus of policymakers and industry to address issues of liability, accountability, and control. Identity controls may be important to help determine accountability and organizational responsibility as well as 'simply' help consumers know where to turn, to which organization, to exercise their consumer and data protection rights.

Private law firms review implications of AI as well as its combination with Distributed Ledger Technology (DLT) and infrastructure. Use-case oriented analysis of combining AI with DLT helps to surface risks and commercial issues not only in legislation specific to these technologies but also other policy and regulatory domains, including data protection, consumer protection, financial services, and private as well as public law considerations. Payment-enabled Agentic AI is one of the topics studied in a joint publication from 2025 by Deutsche Bank and Clifford Chance[6]. They also highlight why cross-chain verifiable identity will be required to enable AI to operate across different platforms.

In a paper on multi-Agent interoperability, Gosmar et al[7] consider technical and operational issues that need to be addressed. They extend analysis of a Multi-Agent model and consider advances that would be needed to *ensure "smooth, efficient, and secure interactions in scenarios where multiple AI Agents need to collaborate".* They posit the role of a Convener that would control access and authority for subsidiary, specialized Agents to interact, highlighting the complex array of delegated authority and scope of intervention that different Agents will need to have assigned to them. They conclude that this will require enhanced specifications to facilitate the framework's security protocols to better protect against unauthorized access and ensure that sensitive information is handled appropriately. This could include implementing advanced encryption methods,

[6] www.cliffordchance.com/content/dam/cliffordchance/Thought_Leadership/convergence-of-ai-and-dlt-final-v2-interactive-compressed.pdf

[7] Gosmar, Diego; Attwater, David; Coin, Emmett; Dahl, Deborah. Ai Multi-Agent Interoperability Extension For Managing Multiparty Conversations. *arXiv:2411.05828v1 [cs.AI] 5 Nov 2024*

robust authentication processes, and more sophisticated mechanisms for managing uninvited Agents.

An IETF draft on AI Agent Authentication and Authorization provides a framework for Agent authentication and authorization and to identify gaps for future standardization. IETF posits that AI Agents should have independent identities and credentials, distinct from users and devices, and explicitly managed workflows. This approach foresees treating Agentic systems as non-human actors with their own identity, credentials, access tokens, and accountability requirements. In standards terms, a program starts to look "Agentic" not just when it reasons, but when it must be treated as an independent actor in authentication, authorization, workflow, and audit design.

Contrasting views are presented from a legal perspective by Chris Reed in his paper on "Autonomy, Responsibility and Agentic AI". He argues that existing legal systems are well equipped to deal with the current situation. The "responsibility gap" created by autonomous Agentic AI is still small, at least from the perspective of criminal law. Legal clarity is most important in the explicit delegation of authority by the user to the AI Agent and recognition of the scope of autonomy accorded the Agent. He emphasizes that contractual rules should normally be sufficient to resolve any questions about the AI Agent's authority and responsibility accorded by the user.

Similar positions about the adequacy of existing law are posited by Mollod et al[8]. They note that terms of service governing the AI Agent are likely to be the first source courts will look to in the event of a dispute. Usage terms provided by vendors and users are likely to shift responsibility as far as possible towards end users. But code law systems may stipulate standard requirements of behavior, transparency or compensation, especially to protect consumers. Scenarios difficult to foresee may however still require consideration beyond the application of existing laws on e-signatures or commerce.

In a 2025 paper by **South, Marro, Hardjono** et al. on ***Authenticated Delegation and Authorized AI Agents***[9] the authors argue that Agentic AI demands "ways to explicitly delegate authority to Agents, transparently identify those Agents as AI and enforce [ ] choices around security and permission". They distinguish three key concepts necessary for defining the scope and requirements of solutions to control of AI Agents: (i) ***authentication*** *confirms an entity's identity;* (ii) ***authorization*** *determines the permissible actions and resource accesses that the authenticated identity is allowed to perform, defining the scope and limitations of delegated activity; and* (iii) ***auditability*** *allows all parties to inspect and*

[8] "Contract Law in the Age of Agentic AI: Who's Really Clicking "Accept"", Cramer, Mollod, Rimer, Rose LLP, New Media Technology Law Blog, accessed on April 3, 2026 at www.natlawreview.com/article/contract-law-age-agentic-ai-whos-really-clicking-accept

[9] South, Tobin; Marro, Samuele; Marro; Hardjono, Thomas et al. Authenticated Delegation and Authorized AI Agents. arXiv:2501.09674v1 [cs.CY] 16 Jan 2025

verify that claims, credentials, and attributes remain unaltered, supporting trustworthy authentication and authorization decisions[10]*.* They review the merits of OpenID Connect and verifiable credentials as means to identify and authenticate AI Agents. OpenID Connect can be complex, requiring multiple sign-in flows to authorize the AI Agent to interact across different service providers. They consider verifiable credential models as an alternative, noting that they can alleviate dependency on central or single provider to mediate trust.

Birch and Hoffart (2025)[11] argue strongly that we will soon need industry wide solutions for Know Your Agent (KYA). They highlight the imminent need for banks to be able to verify the authorizations of Agents to conduct specific activities and the identity and delegation of rights from user to be verifiable. They demonstrate that consumers are already using AI Agents and will continue to drive the need for banks and merchants to adjust to the use of Agents in their fulfillment of both compliance and security obligations. Birch and Hoffart explore the role of digital wallets as tools to manage not only identity but also delegation credentials. They note that, as yet, there is no standardized mechanism to uniquely identify Agents and control their access rights. But they posit that, while the Model Context Protocol (MCP) does not have a framework for authentication and authorization of identities, it could in combination with Policy Decision Points be an enabler for tools that could be used to develop an approach to KYA.

Recently the WE BUILD consortium issued a "non-paper[12]" emphasizing the importance of controlling Agentic AI and highlighting the potential role for the EU Business Wallet architecture. The WE BUILD consortium includes parties such as Bosch, Google Payments, Visa Europe, Ericsson, Spherity. The 'non-paper' proposed to extend the EU Business Wallet architecture trust framework to AI agents to "allow nonhumans to act safely and accountably on behalf of humans, businesses, and public bodies. It notes that "it is important to define the roles that an AI agent can and should play in a digital communication flow of providers, users and relying parties. While this initial paper does not outline a specific architecture for agentic AI credentials or management, it sets the scene for broader debate about how to evolve the baseline framework for business identity and verifiable credentials in the EU to be extended to this use case.

[10] Idem page 1
[11] Birch, David and Hoffart, Jelena. Know your Agent: Enabling autonomous financial services. Journal of Digital Banking Vol. 10, 2 123-134. 2025
[12] 27 February 2026 "*Trusted Identities for AI Agents: An Opportunity for Europe*

3 Policy Concerns and Risks

While most experts agree that Agentic AI can be accommodated within existing principles and laws, there are several policy concerns to which identity solutions may need to respond. These are summarized below.

- *Consumer protection and user confidence*

A primary concern of policymakers is to protect consumers and to ensure minimum levels of confidence in the overall usage of AI Agents. Some common measures may be necessary to educate users, set standards for disclosure and transparency of how Agents operate and ensure that audit trails, access controls and user rights are sufficient to assess liability in the case of disputes. Robust, verifiable identification of Agents and their users will be a prerequisite for clear contractual arrangements and their enforcement. Identity credentials and standards will need to be adapted to the entirely digital nature of AI Agents and integrated into their operational lifecycle and business model.

- *Explainability*

Transparency of data sources and policy rules may be needed to enhance the capacity of users as well as courts to anticipate ex-ante decisions and explain the outcome of an AI Agent's actions. Policies and regulations, an example being the EU AI Act, require that the decisions of an AI tool be explainable and where applicable that users are also informed about how their data may be used in an AI driven process. LLMs may be trained on data from varying sources; users may need to verify the authenticity or accuracy of data on which a model is operating and therefore require mechanisms to trace data access and provenance. AI Agents may therefore need to employ identity credentials that support audit trails of data accessed.[13]

- *Clear allocation of Legal responsibility and liability*

Users of AI Agents will need to assume responsibility for the foreseeable outcomes of delegating tasks to AI Agents. Significant discussion about the legal treatment of AI Agents centers around the need to clarify contractual terms between different actors, including (i) developers and the controllers of, or service providers that use, Agents (e.g., companies that deploy them), (ii) the end users of AI Agents and the companies with which they interact (e.g., merchants), and (iii) the different actors with which AI Agents will interact, including other AI Agents and their controllers. Evidence is needed that parties have understood and accepted often quite complex terms and contractual conditions and taken steps to define or accept policies that define the tasks, scope of autonomy and field

[13] Article 86 of the EU Artificial Intelligence Act introduces a right to explanation for individuals affected by decisions made with high-risk AI systems. Where applicable, users shall be informed about how their data is used.

of application of an AI Agent. Proof of the actions, most importantly non-repudiation of actions taken, and consent of parties will need to be provided in high trust digital media that can be presented and accepted in a court of law if disputes and litigation arise.

- *Data protection and Security*

AI Agents will need access to present payment credentials to regulated service providers or to manage on-chain transactions. This may present a challenge to data security as users will need to share credentials and authorize Agents to act on their behalf. Yet contractual terms for payment services set strict rules on the responsibility and liability of users to safeguard their login and authentication details and to refrain from sharing them with third parties. Agentic AI goes a step beyond open banking regulations that permit third parties to initiate payments on behalf of a client and will require solutions that enable delegation without contravening rules on the sharing of passwords or other sensitive information.

Box 1: Integration with EU Business Wallet Architecture

AI Agent identity and controls may be combined with emerging business identity and wallet ecosystems, such as the framework under development in the European Union[14]. The EU Digital Identity and Business Wallet framework provide a strong foundation for trusted digital interactions. However, complex machine-to-machine use cases such as AI Agents require additional capabilities, for delegation, authority chains, and cryptographically provable mandates.

Identity schemes should consider how to support cryptographically verifiable delegation and organizational roles, without the need to re-issue new credentials for each process or delegation of new or amended powers to an AI Agent. There should be clear mechanisms to trace allocation of responsibility and mandates without disclosure of secure information.

[14] See https://digital-strategy.ec.europa.eu/en/policies/business-wallets

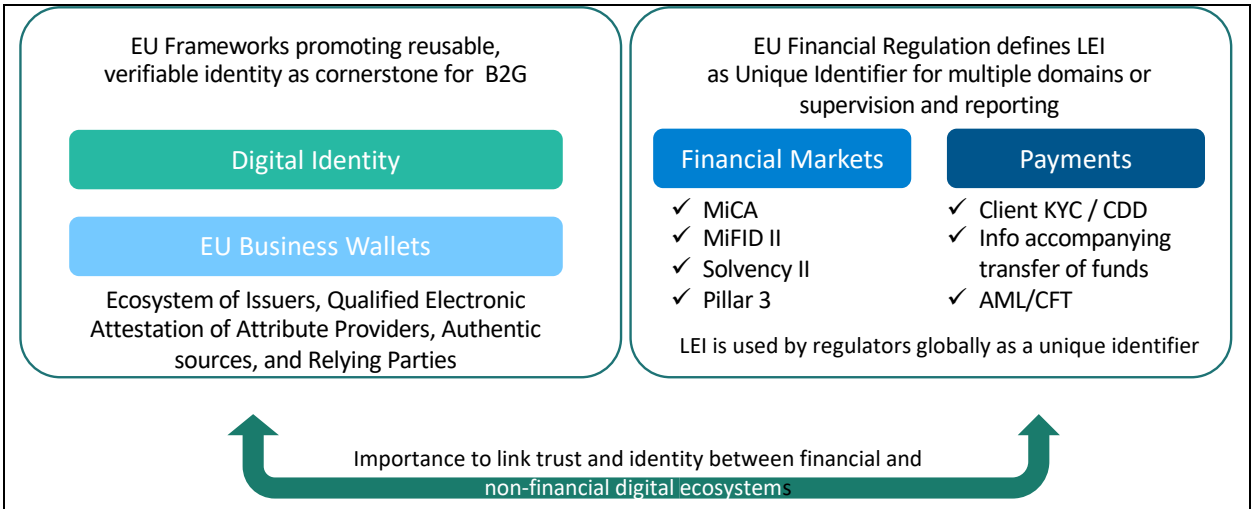

The European Business Wallet architecture should enable multiple identifiers and credentials to ensure global interoperability with organizations outside the EU and support seamless compliance with financial and payment systems reporting and registration requirements set by European authorities. The need to adapt and extend the EU Business Wallet architecture to support agentic AI is already being discussed.
Diagram source: GLEIF

4 Use Case Applications and Requirements

This chapter examines use cases in which Agentic payments create practical value while also exposing the minimum trust and control requirements needed for responsible deployment. The purpose is not to identify one preferred technical stack or business model, but to show how concrete payment scenarios make visible the underlying need for verifiable delegation, machine-readable mandates, runtime policy enforcement, privacy-preserving execution, and auditability across institutional and cross-platform environments.

The preceding chapters argue that Agentic AI in payments should be assessed not only by the efficiency gains it may generate, but also by whether its operation remains attributable, reviewable and consistent with existing legal and compliance expectations. Use cases are therefore relevant to this discussion where they illuminate operational problems that current payment and identity arrangements do not fully solve. They should clarify how an Agent may act on behalf of a person or legal entity without weakening security, obscuring responsibility, or undermining the ability of counterparties and supervisors to verify the basis on which a transaction occurred.

A useful distinction follows from this framing: Some use cases merely demonstrate that an AI Agent can improve user experience or automate a workflow. More relevant for the present discussion are use cases in which the payment itself becomes machine-mediated

and where the economic, contractual or compliance consequences of that payment require explicit proof of authorization and scope of authority. It is these scenarios that most clearly reveal why Agentic payments require new forms of interoperable trust infrastructure.

4.1 Use Cases as a Way to Identify Infrastructure Requirements

Payments have long been highly automated, but existing automation generally operates within tightly bounded institutional systems. User identities, access credentials and approval chains are managed through bilateral or platform-specific arrangements. Agentic AI extends automation into more open, dynamic and machine-native environments. Agents may discover services, interact with counterparties, source data, evaluate options, and initiate payment actions across multiple systems without direct human confirmation at each step.

This shift increases the importance of distinguishing between the identity of the principal, the identity or status of the Agent, the rights delegated to the Agent, and the runtime conditions under which a transaction may proceed. **In many practical settings, the critical issue will be less whether the Agent has a universal or immutable identity, and more whether the parties involved can verify that the specific action taken fell within a valid and reviewable chain of delegated authority.**

The use cases considered below are therefore selected because they expose recurring infrastructure needs. These include: first, verifiable identification of the organization or person on whose behalf the Agent acts; second, a machine-readable representation of the Agent's mandate; third, runtime controls capable of evaluating counterparties, thresholds and policy conditions; fourth, payment execution that does not depend on unsafe sharing of sensitive credentials; and fifth, evidentiary records that support audit, dispute resolution and supervisory review.

4.2 Supplier Payments and Treasury Automation

Supplier payments remain one of the strongest early use cases for Agentic payments because they are repetitive, rules-based, and already subject to substantial internal control requirements. Cross-border supplier payments remain operationally expensive and compliance-intensive despite decades of automation. Invoice validation, confirmation of delivery conditions, checking of vendor status, timing of settlement, liquidity management and foreign exchange considerations often remain fragmented across multiple systems and teams.

In such an environment, an AI payment Agent may create value by assembling and evaluating payment-relevant information, preparing payment intent, and initiating or recommending payment actions within pre-defined limits. The key point is not full

autonomy in the abstract, but bounded automation under clearly defined authority. This is consistent with broader work on Agentic payments, which treats authorization, authenticity and accountability as the foundational challenges once a machine is allowed to transact on behalf of a user or enterprise.

This use case shows that transitioning from conventional automation to Agentic execution will require a more granular trust model. A treasury officer or other authorized person must be able to delegate limited authority to an Agent, specifying the class of payment, relevant counterparties, transaction ceilings, validity period and escalation conditions. Other participants in the workflow must then be able to assess whether that mandate exists, whether it remains valid, and whether the proposed transaction falls within the scope of what was authorized. Depending on the amount, more than one officer or authorized person might also be needed to authorize a payment; in accordance with common risk/compliance requirements, the actual payment instructions usually go through a maker/checker process. Credential management arrangements therefore need to support multi-signature processes that even exist without the application of Agentic AI.

The vLEI-based architecture could be used to create an Agent Mandate Credential (AMC) as well as to encode the authorized mandates that the Agent has been granted. The AMC could be issued by a company representative authorized to do so. The AMC would serve as the 'identity' credential for the Agent. Then the specific functions that the Agent has been authorized to perform, with limits and restrictions, can be included in a separate credential linked to the AMC credential. An example outlined in Annex A depicts how a treasury officer holding an Official Organizational Role credential, creating a cryptographically verifiable chain from the verified legal entity through to the specific Agent. Issuance of Agentic AI AMC credentials, when appropriate, also could be delegated from another type of vLEI role credential referred to in the GLEIF framework as an Engagement Context Role credential (ECR).

The pilot concept in Annex B captures these requirements well. It assumes an operational flow in which an invoice is received, validated against a purchase order, checked against delivery confirmation, assembled into payment intent, presented together with credential references, and evaluated against a runtime policy engine before settlement is executed and a signed audit trail is recorded. This sequencing demonstrates that delegated authority and runtime policy enforcement are related but distinct: the former establishes the right to act; the latter determines whether the specific transaction should proceed at that time and under those conditions. The policy engine evaluates dynamic conditions — live sanctions screening, counterparty exclusion lists, spending thresholds that reset on a periodic basis, jurisdictional compliance requirements — that are not encoded in the credential itself but within whose bounds the credential operates.

The use case also reveals why payment execution for Agents should not rely on direct sharing of user login credentials or other sensitive payment secrets. Existing contractual and regulatory arrangements often assume that account credentials must remain under the control of the account holder or regulated provider. Agentic systems therefore require mechanisms that allow an Agent to prove authorization to initiate a payment without taking custody of the underlying credentials in a way that would be operationally unsafe or legally problematic.

A related implication concerns privacy and confidentiality. In commercial payment settings, transaction details may reveal supplier relationships, treasury strategy, pricing arrangements or commercially sensitive operational patterns. As payment workflows become more machine-mediated, the payment layer should be able to support compliance and verification without forcing unnecessary disclosure of all transaction data to every intermediary or observer. This becomes especially salient in digital asset and stablecoin environments, where machine-native settlement may be attractive for speed or programmability, but where transparent execution environments may create new confidentiality concerns if not carefully designed. Settlement protocols that employ zero-knowledge stealth address techniques can preserve compliance and auditability while preventing unauthorized reconstruction of commercial relationships from on-chain data.

For the purposes of this discussion paper, this use case shows that the value of organizational credentials and privacy-aware payment execution should be understood functionally rather than promotionally. The need is for interoperable proof of organizational identity and delegated authority on the one hand, and for compliant, confidential and machine-compatible payment execution on the other. Whether implemented through existing trust frameworks, new credential structures or specialized settlement protocols, these capabilities respond to operational requirements already visible in treasury and supplier payment workflows.

4.3 Payment for Data in Due Diligence and Compliance Workflows

A second use case of broader significance concerns machine-native payment for data used in due diligence, sanctions screening, beneficial ownership analysis, registry checks, credit assessment and related compliance workflows. This scenario is particularly relevant because Agentic systems increasingly depend on current, authentic and rights-cleared external data, yet many existing commercial models for data access are designed around static subscriptions, manual procurement, or institution-specific integrations that are not well suited to fine-grained machine consumption.

In this setting, an AI Agent acting for a regulated institution may need to engage in several distinct processes. They will need to discover a data provider, determine whether the provider is authentic and authorized to supply the requested information, purchase

access to a signed data package, and retain evidence of both the transaction and the provenance of the data received. The payment is not incidental to the workflow. It is part of the control structure that links acquisition of data to rights of access, cost attribution, and the evidentiary chain needed to support reliance on the output.

HTTP-native payment standards such as the X402 protocol create the technical and commercial foundation for this model. A data provider — whether a credit bureau, national registry, or accounting platform — could wrap an existing data endpoint in a payment-gating layer, specifying the price per query and, where applicable, limiting access to counterparties by presenting a valid vLEI credential chain. An AI Agent authenticates using its delegated vLEI credentials, submits a micropayment at the moment of the query, and receives in return a data package cryptographically signed by the issuing institution and anchored in a KERI transaction event log for provenance. The entire transaction — identity verification, payment authorization, data release, and provenance recording — occurs within a single machine-readable exchange, without human intervention or subscription management overhead.

The commercial rationale is evident and addresses a structural gap in current data markets. Data providers who do not currently monetize their assets — national registries that publish data as a public service obligation, accounting platforms whose users generate verifiable financial records — gain a mechanism to generate per-query revenue without the operational overhead of managing subscriber relationships. Buyers access data from primary authoritative sources rather than aggregators, obtaining stronger provenance guarantees and reducing the risk of relying on stale, aggregated, or potentially tampered data. Birch and Hoffart (2025) identify this as a critical element of the emerging Know Your Agent (KYA) framework: Agents must not only be identifiable but must also verify the identity and authority of the data sources and counterparties with which they interact.

The IFC - FinSpot Digital Invoice Finance Origination (DIFO) protocol provides a concrete deployment context that illustrates the broader applicability of this model. Invoice financing requires lenders to verify the authenticity of invoices, the creditworthiness of buyer and supplier, the absence of sanctions exposure, and the legal standing of transacting entities. Under the DIFO protocol, these verification steps are structured as a sequence of data queries against authoritative sources, each producing a cryptographically verifiable data point that can be assembled into a tamper-evident loan tape. Micropayment-gated APIs for credit data, registry verification, and sanctions screening allow each verification step to be sourced, paid for, and provenance-anchored independently, replacing single-vendor subscription relationships with a composable, auditable data procurement chain. The vLEI credential of the requesting AI Agent establishes that data is being accessed by an authorized party on behalf of a verified legal entity, enabling data providers to restrict access to vLEI-verified counterparties.

This use case is relevant to the discussion paper because it combines several themes developed earlier. These are (i) explainability, (ii) provenance, (iii) machine-to-machine interaction, and the need to verify not only who paid but also (iv) what source of information was relied upon in reaching a decision. It also reinforces the point that Agentic payments are not limited to retail commerce or consumer checkout. They may instead become embedded within business processes where payment functions act as a machine-readable gatekeeper for access to trusted digital resources. Payment, provenance and authorization may therefore need to be designed as part of one integrated control pathway rather than as isolated technical functions.

The infrastructure requirements for this use case are structurally similar to those of the supplier payment scenario, with one additional dimension. The payment environment must preserve confidentiality around the buyer's investigative or compliance behavior — query patterns across credit bureaus, registries and sanctions lists may reveal commercially sensitive intelligence about a financial institution's pipeline — while still allowing legally required controls and records to be maintained. Privacy-preserving transaction infrastructure that prevents unauthorized correlation of individual queries across data providers, while preserving the ability to audit a complete provenance chain for any single credit decision, addresses this requirement without compromising compliance obligations. The pilot concept outlined in Annex C demonstrates the end-to-end architecture for this use case, from vLEI-verified API discovery through X402-compatible micropayment, signed data release, and KERI transaction event log provenance anchoring.

4.4 Cross-Border B2B Settlement and Programmable Execution

A broader application area concerns cross-border B2B payments in which firms seek faster settlement, lower reconciliation costs and programmable execution linked to invoice status, delivery milestones or treasury conditions. This is a natural extension of the supplier-payment scenario, but it introduces additional complexity because multiple jurisdictions, payment systems and compliance rules may be involved. It is therefore an especially useful setting for testing whether Agentic payment models can remain interoperable and policy-compliant when they operate across organizational and technological boundaries.

In this context, digital assets and stablecoins are relevant because they illustrate a settlement environment in which machine-native, time-sensitive and conditional execution becomes technically easier. Real-time or near-real-time settlement can improve automation and reduce frictions, yet it also reduces opportunities for manual intervention and therefore increases the need for ex-ante controls that are machine-readable and enforceable at runtime. Recent analysis of compliance-aware payment architectures for stablecoin rails emphasizes that authorization and compliance logic

should travel with the transaction rather than being left entirely to fragmented ex-post review.

This makes cross-border B2B settlement a useful test case for the control architecture discussed throughout the paper. If an Agent can initiate or coordinate a payment across jurisdictions, counterparties and platforms, the surrounding infrastructure must be able to evaluate whether the relevant corridor is permitted, whether the counterparty is allowed, whether transaction limits are exceeded, whether the mandate remains valid, and whether the organizational identity behind the action can be independently verified. In practice, this means that the Agent's vLEI credential chain must be resolvable by counterparties and payment providers across jurisdictions without reliance on bilateral trust arrangements or platform-specific identity registries.

The wider implication is that Agentic payments require policy-aware execution environments rather than mere payment initiation APIs. The more settlement becomes programmable and always-on, the more important it becomes to embed identity verification, delegated authority, transaction controls and audit logging directly into the payment path. From the perspective of this discussion paper, the significance of this use case lies in how it sharpens the requirement for interoperable trust across different layers. Organizational identity, authority chains, runtime controls and payment execution may be supplied by different actors and technologies. For Agentic payments to scale in cross-border commercial settings, these layers must interact in a way that does not force institutions to rely on brittle credential sharing, opaque platform lock-in or inconsistent evidentiary standards.

4.5 Multi-Agent Workflows and Partitioned Authority

Many practical deployments are likely to involve several specialized Agents rather than one monolithic actor. One Agent may source information, another may verify counterparty status, another may optimize route or timing, and another may trigger the payment itself. This structure may improve efficiency and modularity, but it complicates the questions of responsibility, scope of authority and traceability.

The multi-Agent literature suggests that the real governance challenge lies in partitioned authority and controlled interaction among Agents with different tasks and permissions. A system that treats all machine actions as originating from one undifferentiated actor may be too coarse for meaningful risk management. Instead, there may need to be distinct mandates, rights and evidence trails for the different Agents or components involved in a transaction flow. In the vLEI architecture, this points toward a structure in which an AMC may authorize an Agent to coordinate sub-tasks, but with explicit sub-delegation limits encoded at issuance and verified at each interaction point.

In many operational settings, what matters most is whether the relevant participants can verify what a given Agent is or was allowed to do. It may not be necessary, at least in the near term, to solve universal Agent identity as a stand-alone problem before useful Agentic payment systems can emerge. More important is to be able to identify under whose authority, for what duration, and subject to which controls an Agent has been given delegated controls and powers to act. A focus on rights, mandates and runtime checks may therefore be more practical than seeking an immutable identity model for an ever-evolving software artifact. Companies that use Agents or provide them as a service will need to have tools that equip users with tools to manage delegation and policies that apply to the AI Agents. This could be simply achieved by companies issuing their credentials to Agents, so long as their chain of delegation can be verified. Alternatively, some stakeholders may consider the use of external AI Agent credential issuers that do this on behalf of companies, like trust Agents. This would add additional friction to the system but would align more closely with existing federated systems for trust service providers issuing digital certificates. The economic as well as regulatory benefits and drawbacks of these alternatives would warrant further research.

Use cases involving multiple Agents highlight several additional requirements. Sub-delegation may need explicit limits. Shared event logs or compatible audit records may be required where several Agents contribute to one outcome. Counterparties and payment providers may need standard ways to inspect the authority chain without exposing unnecessary internal information. And where machine-to-machine interactions occur across institutional boundaries, interoperable credential presentation and verification become more important than platform-specific trust assumptions.

4.6 Cross-Cutting Requirements

Taken together, the use cases examined in this chapter generate a common set of infrastructure requirements. These are summarized in the table presented in Box 2 below. They define the design envelope within which solutions for Agentic payment governance must operate and are addressed individually in section 5. No single existing infrastructure element addresses all of them; the cases for vLEI-based delegation models and privacy-preserving payment protocols emerge from use case led requirements rather than prior commitments to a particular technology.

Box 2: Requirements for controlling Agentic AI

Requirement	Description	Relevant infrastructure
Verifiable organizational identity	Agent must be traceable to a verifiable economic actor or legal entity with a globally unique, jurisdiction-agnostic identifier, independent of platform-issued credentials	LEI / vLEI ecosystem; GLEIF public register
Cryptographic delegation chain	Delegated authority from organization to Agent must be machine-verifiable, granular, time-bound, and revocable in real time	Agent Mandate Credential (AMC) / AARC in KERI/ACDC
Separation of credential-time and runtime controls	Static mandate defines scope of authority; a runtime policy engine evaluates dynamic conditions at the moment of execution	AMC defines authorization envelope; policy engine enforces sanctions, limits, counterparty rules
Payment execution without credential sharing	Agent proves authorization to initiate payment without taking custody of account login credentials or other sensitive payment secrets	Delegated wallet infrastructure; credential presentation without credential transfer
Privacy-preserving settlement	Transaction execution on digital asset or stablecoin rails without exposing commercial relationships, volumes or timing to unauthorized observers	Zero-knowledge stealth address protocols; privacy-preserving payment rails
Compliance disclosure on demand	Ability to selectively disclose transaction details to authorized auditors or supervisors without breaking general transaction confidentiality	Selective disclosure; user tainting mechanisms; compliance viewing keys
Data provenance anchoring	Cryptographic record of data sources queried, with signed data packages attributable to known issuing institutions	KERI transaction event log; ACDC verifiable data containers; X402-gated data APIs
Cross-platform interoperability	Agent operates across multiple payment rails, data sources, and institutional systems without platform-specific credential re-issuance	LEI as platform-agnostic identifier; X402 as rail-agnostic payment standard; open credential schemas
Non-repudiation and auditability	Tamper-evident record of Agent actions, mandates exercised, data accessed, available for dispute resolution and supervisory review	Signed transaction event logs anchored in KERI; vLEI credential chain
Sub-delegation controls	In multi-Agent workflows, further delegation of authority must be explicitly bounded, with compatible audit records across Agent components	AMC delegation rules; shared or compatible transaction event log standards

4.7 Implications for the Broader Discussion

Taken together, these use cases suggest that the essential requirements of Agentic payments are becoming clearer even if technical and market standards remain in development. Responsible deployment is likely to require mechanisms that can identify the principal behind an Agent, express and verify delegated authority, evaluate transactions against runtime policy, preserve confidentiality appropriate to commercial settings, and generate records that support accountability and dispute resolution.

This framing helps keep the discussion paper balanced. It avoids presenting any one infrastructure model as the inevitable solution, while still making clear that some functional capabilities will be necessary if Agentic payments are to be trusted in practice. In particular, the use cases show the likely need for interoperable organizational identity and delegated authority frameworks, together with payment execution environments that support privacy and compliance in machine-initiated settings, including those involving digital assets or stablecoins.

Viewed in this way, the relevance of trust frameworks such as vLEI-based delegation models and the relevance of privacy-preserving, compliance-aware payment protocols emerge from the use cases themselves. They are not the starting point of the analysis. Rather, they are examples of the kinds of capabilities that become necessary once one asks how Agentic payments can be made attributable, governable and operationally viable across real institutional workflows.

Box 3: Privacy as a functional requirement

Agent models may need to protect the confidentiality of business data and operations that could be exposed when using public chain ledger-based payments. In commercial payment settings, transaction details can reveal supplier relationships, treasury strategy or pricing; in compliance workflows, the pattern of an institution's own investigations can be revealed. As settlement becomes machine-native and moves onto transparent rails, these patterns become observable to any party watching the chain. Agent-driven volume drastically amplifies the exposure. Privacy is therefore a functional requirement of the payment path in its own right.

The distinction that matters is between verifiability and visibility. A transaction can be cryptographically verifiable, with provable authority and auditable provenance, while remaining shielded from observers who have no need to see it. Techniques such as zero-knowledge stealth addresses, selective disclosure, and predicate proofs, proving a limit is sufficient without revealing its value, allow verifiability and confidentiality to hold at the same time.

The design implication is that privacy-preserving settlement and selective compliance disclosure belong in the same control architecture as identity and delegated authority. The requirement is to disclose to authorized parties on demand while preventing unauthorized reconstruction of commercial relationships from on-chain data. Confidentiality and compliance only end up in conflict when one assumes the only way to verify a transaction is to make it visible. Two examples provide illustrations of the design impact for payments control.

Practitioner example 1: **privacy-preserving transfers and the compliance with zERC20**

- **zERC20 implements private transfers by sending tokens to a one-time, cryptographically derived burn address for which no private key exists;** the recipient later redeems the value by proving in zero knowledge that the address was correctly derived and funded, without revealing the on-chain link between transfer and withdrawal. Compliance controls are risk-based: sending and receiving addresses are screened with blockchain intelligence for exposure to illicit or sanctioned activity. The example shows both sides of the tension this section describes: transactional privacy is achievable with proof-based mechanics, while today's controls attach to addresses and their history rather than to the authority under which a transfer is made. Anchoring such mechanisms to governed identity and mandate evidence, with disclosure to authorized parties rather than to all, is the complement this paper's trust layer describes.

  In an upcoming release, zERC20 will be supporting account-specific regulatory disclosure through dedicated "view keys". Upon receiving a valid and legally binding request, the operator will be able to provide an authorized investigative authority

with a view key for the relevant account without requiring the account holder's consent. The key enables the authority to reconstruct that account's incoming and outgoing transaction history, while preserving the confidentiality of unrelated accounts and transactions.

Practitioner example 2: identity-anchored confidential payments using Curvy Protocol

- Curvy implements confidentiality through one-time stealth accounts and zero-knowledge aggregation: each incoming transfer lands on a fresh, unlinkable account, amounts are shielded, and deposits are separated from withdrawals so that flows cannot be correlated on the public ledger. Compliance attaches at the identity level rather than the transaction level: a Curvy identity carries screening and policy logic, operators hold viewing keys that make their own flows fully auditable end to end, and selective disclosure shares specific facts with specific counterparties without opening the account's full history. The same identity primitive extends to AI Agents: an Agent receives a payment identity with counterparty rules and operator-held viewing keys, compatible with emerging Agent payment rails. The example illustrates the direction this section describes: confidentiality by default, with control and disclosure anchored to programmable identities rather than to ledger transparency. Binding such identities to governed legal-entity evidence is the step this paper's trust layer describes.

5 Operational Challenges: From Principles to Practice

Several practical issues arise as stakeholders attempt to apply existing identity and control frameworks to AI Agents in payments. This section outlines operational challenges, requirements and design considerations in the context of using Agentic AI in payments operations. Identity credentials will need to support operations in which an Agent is able to establish proof of access, mandates or actions in ways that prove what actions were taken, under what authority and ensure security and auditability along the chain, compliant within existing legal frameworks. This section draws on insights and positions outlined in the literature review and shares considerations based on GLEIF's role in the provision of high-trust, digital organizational identity solutions.

- **The composition and lifecycle of AI Agents make their identification complicated**
  AI Agents are produced by vendors, evolve with the enhancement and training of their models and are composed of multiple tools and components. They adapt over time and may have different capabilities and decision-making results that require dynamic controls to be put in place rather than static identities. The "identity" of an AI Agent can be tied or change based on its underlying model, its orchestration logic, the tools it can access, or the policy and mandate under which it operates; these attributes of the Agent may all be relevant for how its identity or role is defined. This definition is critical, as it directly affects how respective credentials are issued, managed, and revoked.

- **Users require granular, role specific digital identity**
  A minimum requirement for responsible AI usage is the ability to clearly identify the user of an Agent and trace their activity. Identity credentials are needed to accept Agent's contractual terms, delegate powers to the Agent and confirm the terms of policies that govern the autonomy and scope of authority accorded to the Agent. Identity needs to be granular and often specific to the context of transactions and the roles of individuals acting on their own behalf or on behalf of a legal entity. Consumers and businesses will frequently use even the same AI Agent in different roles, applying different policy rights and obligations. The role of an Agent for instance by a compliance officer may be different from the tasks allocated to it or an array of Agents to perform cash management functions such as to optimize settlement reserves or conduct sweeps.

- **Mechanisms are needed to issue and manage delegation of authority to Agents**
  Agents need to present their authorization to the other parties, including other AI Agents, with which they interact. Delegated authority needs to come from users or the service providers that engage the use of Agents on their behalf to fulfill business operations. This chain of delegated authority needs to be verifiable and anchored in credentials that can be issued to Agents, defining the specific rights they can exercise in a way that third parties can process. Delegated credentials need to be accessible for users to easily and reliably manage, revoke and adjust in real-time without over-reliance

on third parties such as certificate authorities to re-issue or alter existing credentials each time.[15].

- **Identity of Agents may be manifest in the rights they have and actions they perform** An open issue is that Agents are difficult to define let alone assign a unique, independent identity. Clear, unambiguous definitions of an Agent may be difficult to pin down. Agents refer to an assembly of interfaces, LLMs, data sources on which it can be trained and policies governing its usage. Agents may regroup multiple sub-Agents or tools. The status and usage of an Agent may evolve over time. While there are principle-based frameworks (see the IETF working proposal), practical solutions may take time to develop into robust standards. In the meantime, assignment of a definitive and immutable identity to such artefacts may remain problematic.

  An alternative proposed by practitioners is to focus on the management of rights and control of actions performed by Agents, rather than immutable control of identity. This perspective argues that it is sufficient to track the actions of an Agent and confirm at each step that their powers have been duly delegated and their roles aligned with the task they undertake. For example, an Agent may be accorded rights to an address on a ledger to enact transfers under certain conditions set out in the policy framework. These delegated powers could be then controlled and verified against the rights and identity of the delegating party (e.g., the User).

  Hardjono framed the problem through a multi-actor liability lens. In a typical Agentic payment chain there are five distinct entities, each with separate legal and financial obligations: the consumer, the AI provider, the cloud provider where the Agent executes, the payment startup that exposes the rail, and the bank that holds the underlying account. The obligations of each remain unclear by default when the action is initiated by an autonomous software Agent. He pointed toward vLEI as the appropriate primitive for anchoring this complexity, because the credential binds the Agent's activity back to a legal entity. The static elements of the identity stay anchored. The dynamic record of how the Agent has changed accumulates around that anchor without replacing it.

  Marija Vlajković framed a complementary distinction from the regulatory side: static authorities versus dynamic policies. Authorities of an organization, of an officer holding a role, of an Agent acting under delegation, are reasonably stable and are appropriately encoded in credentials at issuance. Policies that govern whether a specific transaction may proceed at a specific moment are inherently dynamic, evolving with sanctions updates, counterparty status changes, budget decisions, and jurisdictional rules. Encoding all of them in the credential at issuance is neither operationally tractable nor

[15] Refer to South, Marro and Hardjono in their assessment of trade-offs using OpenID Connect

architecturally desirable. The mapping onto Hardjono's technical framing is direct. Credential-time captures static authorities. Runtime captures dynamic policies.

- **Machine-to-Machine authentication requires identity credentials for Agents**
  Credential frameworks will need to be adapted to technical and operational environment without a "person in the middle". While traditional processes can rely on persons to authenticate and validate a transaction that is prompted by a machine, and therefore look out for anomalies and assert control, machine-to-machine transactions will require more automated, cryptographically secure ways to confirm their role and authority to act on behalf of their users. Traditional user login identifiers, customer records and one- or two-factor confirmation methods may in an online environment be impractical to transpose to Agents. Detailed design issues warrant further consideration to evaluate the adequacy of alternative credentials.

- **Payment credential management needs to comply with data secrecy rules and contractual terms**
  Agents should be able to authenticate and authorize transactions without contravening rules set by payment service providers on safeguarding account login and authorization credentials or other data subject to strict data security rules. The storage, safe-keeping and regular update of multiple payment service credentials with AI Agents may pose not only operational and security issues. Contractual terms with banks or other payment service providers may explicitly prohibit the sharing of credentials with third parties. The role of outside developers, Agents managed by third parties and even Agents delegated by users, may require data security rules to be revisited. Different technical mechanisms such as tokenization and privacy enhancing verifiable credentials may be more appropriate mechanisms to use to authenticate and authorize transactions initiated by an Agent.

- **Cross-platform and DLT operability require management of multiple platform-generated identities**
  Cross-platform operations are challenged by proliferation of multiple identity credentials issued by platforms to users. Transaction verifications, such as of the signatory's rights and limits or of the payee account will require verification of account controller's or owner's identities; And actual instruction of payments and the authorization of those transactions will require account or payment service providers to authenticate and ensure that transactions are authorized by the appropriate parties. This could require sharing multiple platform login and security credentials with AI Agents. Changes to the credentials managed by each platform, payment system or service provider would complicate the operational, legal and security arrangements for these Agents.

- **Data provenance and transparency may be needed to ensure 'explainability'**
  AI Agents rely on various data sources to train LLMs and make proposals and decisions. Means to assess and demonstrate data relevance, authenticity and accuracy may be necessary if outcomes are to remain 'explainable' and defensible. Identity credentials can help to trace access to and provenance of data within these ecosystems. Identity solutions for AI should consider not only how to be checked when sourcing data; the identity of data sources or controllers may need to be verified to ensure that incorrect, unauthorized or erroneous data is not used in contravention of an agreed policy.

- **The role of Wallets and verifiable credentials**
  Experts also note that many Agent-driven tasks will need to commit persons or entities to contracts in which the identity of an end-user needs to be explicitly verified or used to assign rights or assets. A simple example in consumer services might be the purchase of airline tickets; a corporate use case might be the submission of official reports or declarations of a company officer. The AI Agent may need to have access to a digital identity wallet or credential and be able to 'present' this within the overall process, without manual recourse to a human to finalize the operation.

16

Figure 2: Composite agentic identity ecosystem actors

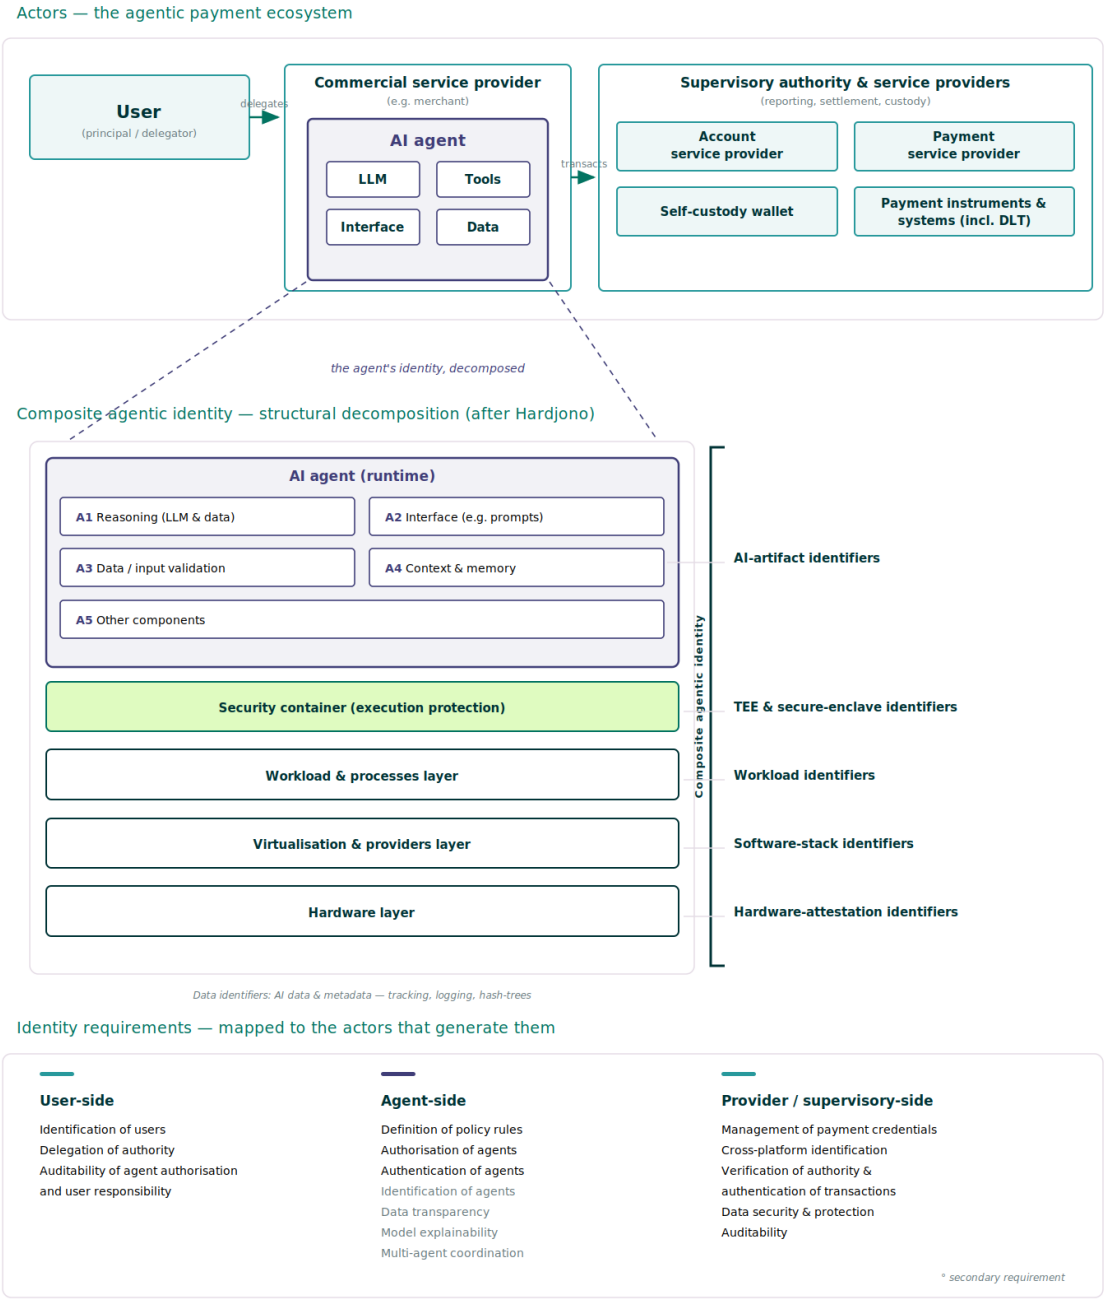

Source: GLEIF

[16] Structural decomposition of the Agent's identity (after Hardjono), and identity requirements mapped to the actors that generate them (Authors' own illustration)

6 Design Principles, Options and Open Issues

Architectural design depends on the context in which AI Agents are applied. This section outlines considerations related to credential management and standards as well as the underlying structure of AI Agents and their economics. An AI Agent is a piece of running software; it may not be immediately clear who holds legal and financial obligations for different parts of it and its usage. Like in other areas of technology, there are multiple parties involved. Moreover, the Agent is not a monolith – it is an assembly of multiple components, including usually sub-Agents, which by design can evolve over time through new training events and access to different data sets, tools and capabilities.

The workshop[17] surfaced a further set of design issues and viewpoints that this paper engages with but does not, on its own, fully resolve. Each is stated below together with the bearing it has on a delegation-based approach to organizational identity and machine-readable mandates:

1. **Architectural approach and standards coordination.** There is a proliferation of standards and approaches that now need coordination. There are concurrent and potentially very complementary efforts spanning (i) Agent authentication and authorization frameworks, (ii) Agent-identity and registry protocols, (iii) hardware-anchored identity registries and (iv) the binding of verified organizational identity into asset-transfer protocols. An emerging position in this paper is that standards development should focus more on binding authority over Agentic AI to existing organizational identity frameworks, rather than the creation of new identities specific to AI. There are robust organizational identity frameworks to which Agentic AI can be anchored to express delegated authority as a machine-readable, revocable mandate. But further analysis and more refined frameworks are necessary to expose where there are complementarities and flag where coordination – for instance on credential-presentation and verification interfaces at the payment instruction layer - are preconditions for efforts to interoperate as well as to support further delegation of Agent AI credentials by AI Agents themselves.

2. **Issuer models for the attribute layer.** Control over an Agent's deployment context, training state, and mandate updates must originate somewhere, and the choice of issuer has consequences for both trust and verification cost. Three models are plausible: (i) a cloud or infrastructure provider acting as a direct issuer; (ii) the buyer organization (of the Agent) issuing to its own Agents while drawing on provider metadata; or (iii) a third-party attestation service issuing and managing these credentials on the organization's behalf. The model proposed here favors (model ii) issuance anchored to a verified legal entity, so that whichever party issues attribute

[17] This refers to the hybrid workshop convened in April by GLEIF and Curvy.

credentials, the chain resolves back to an identifiable organization rather than to a platform-specific account. The buyer-as-issuer model aligns most directly with that principle, since the delegating entity is also the credential source; provider-issued or third party-attested models remain viable but introduce an additional party whose own identity and authority must themselves be verifiable. The trade-offs across these models are primarily economic and are examined in Annex D.

3. **Credential form for the Agent**. A related and still-open choice concerns the credential form assigned to the Agent itself: whether to reuse an existing Engagement Context Role[18] credential, define a new Agent-specific credential within the existing trust hierarchy, or treat the Agent as a delegated identifier established under an existing human identity. The architecture in Annex A leans toward the third option; the Agent as a delegated identifier under the issuing human's key event log because it binds the Agent to a principal from inception without asserting that the Agent occupies an organizational role in the way a person does. This choice interacts with a deeper question: because an Agent is continually retrained and recomposed, it may be better understood as an evolving actor than as a fixed identity. That argues for keeping the core identifier stable while allowing the attribute layer to absorb change, consistent with the credential-time/runtime separation set out above, and for recording the validity window over which a given configuration of the Agent was authorized to act.

Box 4: Flexible attribute layer management design options

Option (a): Edge-chained credentials. Each attribute event is its own ACDC, chained back to the Agent's identity ACDC through the Edge section of the credential. Training events, deployment context updates, mandate updates, and transaction authorizations would each be issued as separate credentials in the chain. The advantages of this model are that it can be natively supported by the ACDC format and that it is infinitely extensible, with no schema changes required at any point. The downside of this approach is that chain traversal cost at verification time grows with chain depth, and the operational complexity of maintaining many credentials per Agent over its lifetime is non-trivial.

Option (b): Schema variants. Possible attribute variants are pre-baked into the credential schema at issuance, and the credential carries one of N pre-committed forms. The benefit of this model is that it is a single credential, no chain traversal at verification. The downside to this design is that future variants cannot be anticipated at issuance time, the schema bloat grows combinatorially

[18] This is the official term under the GLEIF vLEI Ecosystem Governance Framework for a credential that can be issued by someone holding an official organizational role in an entity to any staff, partner or other party with a specifically defined (i.e. context specific) role to engage with or on behalf of the entity.

with attribute types, and migration is required whenever a new variant type emerges in the ecosystem.

Option (c): External attribute logs anchored by hash. The credential commits to the hash of an external log, which may be a database, an event stream, or any other tamper-evident structure. The Agent reveals log entries on demand and the verifier checks them against the commitment in the credential. The advantage of this approach is that it is flexible and will not bloat the credential, nor is chain traversal required at the moment of verification. The disadvantages of this model are that it requires trusted external infrastructure, and that the verifiability of the external log is weaker than via a native ACDC mechanism unless additional protocols are added.

4. **Attribute layer management and evolution**. A flexible attribute layer is needed for the context of Agentic AI. Different models and mechanisms should be considered based on the level of assurance, market framework and economics. The box below considers three technical options based on use of Authentic Chained Data Credentials (ACDCs)[19]. Each option has its use-case advantages: Edge-chained credentials may fit use-cases for institutional B2B transactions where chain depth remains bounded, and verification cost is not an impediment. Schema variants are better adapted to consumer-rail scenarios where the set of attribute types is well known in advance. External attribute logs address high-frequency machine-to-machine scenarios where verification economics are dominant. Annex D sets out why the verification cost of these options diverges sharply with transaction volume.

5. **Policy-engine operator.** Because credential-time and runtime are distinct functions, the party that operates the runtime policy engine need not be the same party that issues the credential. The choice has commercial and regulatory implications. Candidate operators include the buyer's bank, the settlement rail, and the organizational-identity layer itself, each optimizing for different priorities of regulatory recognition, latency, privacy, and neutrality (these are examined in Annex A). The framework here is deliberately agnostic on this point: it specifies how authority is delegated and made verifiable, not who must enforce dynamic policy. What it does require is that whoever operates the policy engine can resolve and verify the delegation chain presented to it, so that runtime enforcement and credential-time authority remain consistent regardless of operator.

[19] While this format is only one specific example that could be used, it is the basis for chained credentials that underpin the decentralized vLEI organizational identity and allows for autonomous issuance, revocation and management of credentials adapted to the needs of complex network-based analysis. These examples are provided for illustrative purposes and discussion.

Box 5: Practitioner example of decoupling reasoning and execution

Decoupled reasoning and execution as a runtime enforcement architecture - example

One approach to **the runtime policy-engine question**, developed by Terminal 3, an agentic security platform working with financial institutions and digital identity authorities, is to architecturally separate an Agent's *reasoning* from its *execution*.

Under this model, planning may be performed by any AI model, which is treated as untrusted; execution is performed by a separate, hardware-secured runtime (using trusted execution environments) that holds credentials and secrets, enforces machine-readable policy deterministically at the moment of action, and produces a tamper-evident audit record of every action taken. Because the execution layer - not the model - presents credentials and signs transactions, payment and account credentials are never exposed to the reasoning layer, addressing the data-secrecy and credential-custody concerns raised in Section 5.

The approach is consistent with the credential-time/runtime separation set out in this paper: verifiable organizational identity and delegated authority (for example, vLEI-anchored mandates) are established at credential-time, while the hardware-anchored execution layer enforces dynamic policy at runtime. This architecture is designed for contexts including KYC-verified whitelists for regulated stablecoin distribution and national digital-identity integrations for Agent onboarding.

6. **Credential vs. Run-time.** Designs should reflect the difference between credential-time and runtime functions. The credential carries the authorization directly. Scope, identity, and delegated authority live in the credential at issuance. But many if not most of the policy-related decisions that matter in Agentic B2B payments will happen later, based on the exercise of that policy[20]. The runtime policy engine determines whether the specific action proposed at the moment of execution should proceed. Credential-time needs to capture static authorities. Runtime must capture dynamic policies. The mandate itself divides along the same line. Its document (who may act, within which scope and limits) is fixed at issuance and can be cached. Its state (how much of a limit is already consumed, whether the delegation has been revoked, whether each link in the chain is still live) can change between any two actions. Card authorization routinely verifies dynamic state in the form of available funds at the

[20] For instance, a spending threshold might be breached; a sanctioned body could be added to an external record; a counterparty status could change during an overall transaction; or a mandate could be updated by the principal. None of these live in the credential at the moment at which it was issued. They relate instead to external events. The credential anchors the policy rights, defining who is allowed to act and within what scope.

moment of the transaction; for Agent-initiated payments the state that matters is the authority itself, and on open rails no single intermediary holds it.

7. **Stablecoin rails and the policy engine.** A specific instance of the operator question is whether the assumption of a regulated balance-sheet intermediary still holds when settlement occurs on a stablecoin rail or a privacy-aware rail rather than through a bank. The delegation model is rail-independent by design: the credential chain authorizing an Agent to act is the same whether settlement is on a conventional rail or a programmable one. What changes on these rails is where compliance logic must sit. Where there is no regulated intermediary in the settlement path to perform sanctions, AML, and mandate checks, that logic must travel with the transaction or be enforced by a policy engine at the point of execution. This reinforces the case for embedding verifiable authority and runtime controls into the payment path rather than relying on an intermediary to supply them after the fact.

8. **Human accountability and autonomous execution**. A genuine tension exists between two directions of travel. In many jurisdictions, machines cannot act as legal agents for business operations, and a responsible human must remain in the chain of accountability. At the same time, recent industry protocols have introduced explicit support for autonomous execution in which Agents transact on pre-authorized instructions without human supervision of each individual transaction. These positions are not necessarily incompatible: a pre-authorization can satisfy the requirement that a human stands in the chain of accountability even when no human is present at the moment of execution. A delegation model anchored to a verified individual within a verified organization is well suited to expressing exactly this — the human's authority, scope, and time bounds are captured in the mandate at issuance, and every subsequent autonomous action resolves back to that human-anchored authorization. The residual question is jurisdiction-specific: whether a given legal system treats a bounded pre-authorization as sufficient human agency. The model can record and prove the pre-authorization; it cannot pre-empt how a court will interpret it.

9. **Cloud-hosted versus self-hosted deployment.** The paper should be explicit about whether it assumes cloud-hosted deployment, self-hosted deployment, or both, because issuer models, policy-engine operator choice, attribute-layer governance, and verification economics all behave differently across that boundary. Self-hosted deployment fits institutional B2B settings with mature security infrastructure, where an organization controls its own key management and issuance. Cloud-hosted deployment fits faster iteration, smaller organizations, and consumer-adjacent use cases, but shifts custody of keys and credentials toward the host. The delegation framework itself is deployment-neutral, the chain of authority is the same in both cases but the question of who holds the Agent's signing keys, and therefore who can be compelled or compromised, differs sharply. The paper should advocate clearly for

the institutional, self-hosted profile in the B2B scenarios it treats as primary, while noting that cloud-hosted deployment is a legitimate profile for lower-assurance contexts.

7 Conclusions, Requirements and Approaches

This paper has considered the challenges posed by using AI Agents in payment services, outlined requirements they set for identity infrastructure and introduced different approaches to manage delegated authority. AI Agents will be accorded increasing autonomy and interact in an ever more automated, digital environment. This will heighten the need for granular, cryptographically secure mechanisms not only to verify identity but more importantly the chain of delegation of rights and authority accorded to Agents. Verifiable credentials and wallet infrastructures may present important foundations to move away from low assurance, platform-issued and controlled identity to high trust credentials that can operate in both legacy and DLT ecosystems.

A focus must be on what new practical tools are needed to address existing policy and regulation. Policy makers and industry players tend to agree that while envisioned models of deployment of AI-Agents do not fundamentally challenge existing principles and laws, new means may be necessary to translate them into practice. Assuring transparency, accountability and fairness will require digital identity and delegation tools that fulfill the need for

- High security
- On- and off-chain viability
- Scalable delegation of authority and management of policies
- Alignment with existing financial sector regulation and reporting infrastructure
- Cross-boundary (jurisdictional, organizational) validity and trust
- Time-stamped, traceable auditability and 'explainability'

Tight control of the actions of AI Agents may compensate for lack of universal means to identify them. The paper has noted that further discussion is needed to move from principles to practical measures to assign unique identifiers or identities to Agents. AI Agents are compositions of multiple software elements that due to their inherent capacity to 'learn' are also in evolution. They run on hardware systems that may also be a source of identity and control. Multi-Agent systems are also prevalent in the current landscape. Users are likely to interact with many different programs, data sets and policies. Pinning down a precise, immutable and cross-industry identity standard for an Agent may be a quite complicated and lengthy process.

In the meantime, delegation of rights by a User or Provider may be sufficient to fulfill necessary operational or legal objectives. Controlling not so much 'who' Agents are, but who controls them and what they are allowed to do, and under what authority, may be a more practical and pertinent incremental approach. This can be achieved in parallel with identity standards by harnessing the capacity of systems like the vLEI to already issue granular, policy and time-bound credentials for delegation of authority.

Avenues for further research and development.

This paper and discussions contributing to it have also raised further practical operational and policy questions that warrant further investigation and analysis. A non-exhaustive list of issues noted by the authors is shared below.

- **Assurance levels for mandates.** Identity assurance frameworks grade how a subject was verified. For Agent-initiated transactions the evidence relied upon includes the mandate itself. Further work should define assurance levels for delegated authority, covering how the delegation was verified, the precision of its scope and the freshness of its revocation status, and examine how liability allocation can follow mandate assurance in the way it follows authentication strength in card schemes and signature levels in electronic-signature law.

- **Verification economics at machine cadence**. Annex D outlines how verification cost scales with delegation depth and volume, and how caching bounds the freshness of revocation data. Further work should treat acceptable staleness as an explicit design parameter: how fresh authority evidence must be for a given action rate, settlement finality and strength of recourse; who sets that bound; who bears the residual risk inside the staleness window; and how amortization behaves where open discovery makes first encounters dominate. Proof-based approaches that compress verification to near-constant cost should be evaluated against the same parameters.

- **The reliance record and recourse.** When a relied-upon check later proves wrong, recourse depends on evidence of what was verified, by whom and under which conditions at the time of reliance. Further work should specify the minimal transaction-time record that makes reliance auditable and disputes resolvable across parties, without turning the record itself into a surveillance artifact.

- **Binding authority to Agent state.** Section 6 treats the Agent as an evolving actor whose attribute layer absorbs change, with mandates valid for a given configuration. Attesting the running state remains open: establishing that the deployed Agent is - at the moment of action – engaging and acting based on the configuration that was authorized. Bridging credential-based authority with runtime attestation of software state deserves dedicated work.

- **Integration into national business identity and wallet schemes:** with a growing number of national frameworks for business identity emerging, work is required to see how their architectures may align with or support the complex and shifting requirements for Agentic AI. Business Wallets should plan to integrate credential ecosystems and identifiers that help participants meet existing reporting and legal requirements while also supporting future, more advanced requirements of AI-driven processes. For

instance, this means enabling platform independent auditability, chaining of delegated authority and the ability for cryptographically secure signing and verification. National schemes should anticipate the use of AI-Agents that operate on a machine-to-machine basis, span multiple jurisdictions and require both ex-ante and ex-post reporting to supervisory and regulatory authorities to comply with different payment service-related obligations.

Annex A: vLEI-Based Agent Mandate Framework

The vLEI ecosystem provides a practical foundation for establishing cryptographically provable mandates for AI Agents. This annex outlines a proposed credential architecture, the Agent Mandate Credential (AMC), that extends the existing vLEI trust hierarchy to cover delegation of authority to AI Agents in payment contexts.

The existing vLEI credential chain links a Legal Entity Identifier to specifically identified persons within an organization. This can enable verification of authority and role based on either an Official Organizational Role (OOR)[21] and/or an Engagement Context Role (ECR) credential, which can be defined autonomously by a legal entity but using a scheme agreed upon with other parties. The AMC would sit below this layer, issued by an authorized human credential holder to an AI Agent's cryptographic identifier. Rather than identifying the Agent as a person in a role, the AMC encodes the scope of what the Agent is permitted to do: the class of actions it may take, any transaction limits, permitted counterparties, and the time window within which the mandate is valid. Crucially, the Agent's identifier would be established as a delegated identifier under the issuing human's key event log, cryptographically binding the Agent to its principal from the moment of inception and creating an unbroken chain of accountability back to a verified individual within a verified organization.

In a practical payment context this would work as follows. A CFO holding an OOR credential issues an AMC to a treasury Agent, authorizing it to execute FX sweeps up to a defined transaction limit within a specified validity period. When the AI Agent identifies a sweep opportunity and constructs a payment instruction, it assembles a credential presentation comprising the full chain from LEI through OOR to AMC and submits this alongside the instruction to the Payment Service Provider (PSP). The PSP resolves the issuer's identity, verifies the credential chain, confirms the mandate has not been revoked, and validates that the transaction falls within the authorized scope before executing. The PSP retains its own payment rail credentials internally, the Agent's vLEI chain authorizes the right to act, it does not expose sensitive account credentials. Each action taken by the Agent is recorded as a signed event, producing a tamper-evident audit trail that supports both operational monitoring and post-hoc dispute resolution.

[21] GLEIF references the ISO standard 5009 for official roles (see https://www.iso.org/standard/80603.html) and incorporates them into the credential standards for vLEI OOR credentials (see https://www.gleif.org/en/lei-data/code-lists/iso-5009-official-organizational-roles-code-list)

Figure 3: Architecture for an Agent Mandate Credential

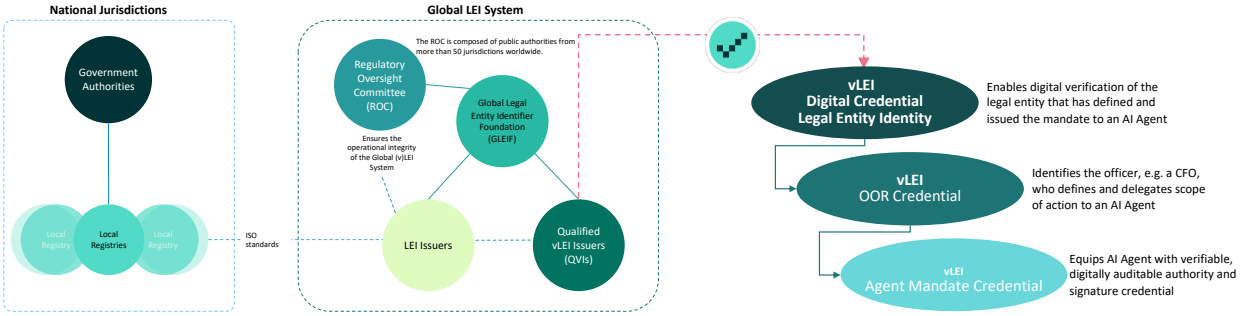

Sources: GLEIF

This approach addresses several of the operational challenges outlined in the main paper. It provides granular, revocable delegation without requiring the Agent to hold or present sensitive payment credentials directly. It produces the cryptographically provable chain of intent and authority that regulators and courts are likely to require. And it does so by extending existing vLEI infrastructure rather than introducing entirely new identity primitives.

The linkage of an AMC into the GLEIF ecosystem enables interoperable audit and control. Authority mandates issued via the chain of vLEI Credentials enable any counterparty to trace back the chain of issuance and delegation to any corporate or legal entity using public high assurance level 1 data accessible via GLEIF. The GLEIF ecosystem is mandated to ensure interoperability of legal entity identities and enable both public and private entities to verify who they are doing business with and how entities and persons acting on their behalf are linked via the chain of credential issuance without operational or legal reliance on a third party to attest to their relationships. This makes complex network analytics across different jurisdictions and legal entities more scalable without sacrificing the ability to audit decision making trails.

Several questions warrant further development by an industry working group. The formal ACDC schema for the AMC needs to be defined, including the precise field structure and validation rules. The PSP-side verification API represents the critical adoption dependency, without a standard interface for credential presentation and chain verification at the payment instruction layer, the model cannot scale. The question of multi-Agent delegation also requires consideration: where an Agent spawns or coordinates with sub-Agents, it is not yet clear whether the AMC should permit further delegation or whether each Agent in a chain requires its own independently issued credential. Finally, the relationship between this framework and emerging protocols such as Google's Agent Payments Protocol and the IETF Agent authentication draft should be explored to identify opportunities for interoperability rather than parallel and competing approaches.

- **Run-time vs credential time:** The policy engine that evaluates whether a specific transaction should proceed at runtime is a distinct architectural function from the credential that authorizes the Agent to act. As the Section 6 supplement sets out, credential-time and runtime are related but separate. The credential anchors who is authorized and within what envelope. The runtime policy engine determines whether the specific action proposed should proceed at the moment of execution. This raises a structural question: who operates the policy engine in a multi-party payment scenario, and what are the commercial and regulatory consequences of that choice?

- **Positioning policy management with the appropriate entity:** Five candidate operators are emerging from current market activity, each optimizing for different priorities.

    - **The buyer's bank as policy engine operator.** This option is regulated, recognized, and fits the existing legal liability chain. The bank is already a counterparty to most of the policy decisions that matter, including sanctions, AML, and internal mandates. Banks have been slow historically to ship Agent infrastructure, and most lack scale economics at SME level, though several major institutions including JPM Kinexys are now actively experimenting in this space.

    - **The payment rail itself, or a privacy-aware settlement layer, as policy engine operator.** This could enable fast, privacy-native policy management to be directly embedded in the payment rails. But the regulatory perimeter is unclear for cross-border activity. This option concentrates trust in the rail provider, which has its own risk profile.

    - **The identity-attribute layer as policy operator:** an identity network provider or issuer network (such as GLEIF or a QVI) could fulfill this role of policy engine operator. The advantage is that it would provide a neutral, standards-based arrangement that avoids locking buyers to a specific bank or rail. However, this option does not yet exist at operational scale and the actual governance, adaptability, responsiveness to commercial needs and funding models remains open.

    - **A specialized payment startup acting as policy engine operator:** this could enable policy to be managed independently of the payment rails. This model is implicit in several existing agentic commerce offerings. It optimizes for speed of implementation and offers a clear commercial point of accountability, but it duplicates several functions performed by other entities in the chain.

    - **The corporate self-hosted policy engin**e: The buyer organization can run the policy engine inside its own perimeter, integrated with its existing ERP, treasury,

and compliance systems. This option gives the largest measure of control over data and decision logic. It does not scale to SMEs. Recent infrastructure announcements have made this option architecturally realistic in a way it was not at the time of the workshop, as set out below.

Commercial and regulatory incentives will shape the choice of operator as much as architectural fit. Each candidate has reasons to want the policy engine to sit with them. It is preferable to be explicit about that tension than present one model as the natural conclusion. Different use cases will likely require different operators. Cross-border B2B settlement, the consumer-rail card scenarios, Agent-to-Agent data marketplaces, and treasury automation each impose different priorities on the policy engine.

The corporate self-hosted operator scenario was theoretical at the time of the workshop and is now architecturally realistic. The scenario depends on infrastructure that allows Agent execution and tool calls to occur inside the buyer's perimeter while the Agent loop itself remains hosted elsewhere. On 19 May 2026, Anthropic announced two infrastructure capabilities for Claude Managed Agents that change the practical realism of this scenario. Self-hosted sandboxes keep the Agent loop on Anthropic-managed infrastructure while running the Agent's tool execution inside the environment controlled by the buyer organization, either via its own infrastructure or that managed by a provider such as Cloudflare, Daytona, Modal, or Vercel. MCP tunnels allow the Agent to reach Model Context Protocol servers inside the buyer organization's private network through a lightweight outbound gateway, without exposing those servers to the public internet.

With self-hosted sandboxes, a further split of roles between the buyer organization and the AI provider can be considered. The Agent's tool execution, the policy engine, and the integration with internal systems could all be run within its own security perimeter, while model inference would remain on the AI provider's infrastructure. This split is more comfortable for many large institutional buyers than either pure SaaS or pure self-hosted, because it preserves data sovereignty and audit trails inside the perimeter while leaving the operationally complex model serving to the provider. For mid-market organizations without mature security infrastructure, the trade-offs likely still favor SaaS deployment. The same capabilities bear directly on issuer model (b) described in Section 6 and the Annex A supplement, where the buyer organization issues attribute layer credentials drawing on internal metadata APIs: MCP tunnels give the Agent a concrete mechanism to reach those APIs without exposing them publicly.

Disclaimer. The Anthropic infrastructure capabilities referenced above are factual and verifiable as of 19 May 2026. The architectural interpretation of how they intersect with the vLEI operator and

issuer model questions is interpretive and should be reviewed by the GLEIF technical team, and by Anthropic, before publication if the paper intends to reference this specific announcement.

Box 6: The GLEIF ecosystem and application to identification and control of AI Agents

The Legal Entity Identifier (LEI) is a non-proprietary globally recognized identifier for legal entities. It uniquely identifies organizations participating in financial and digital ecosystems. The LEI is recognized within the eIDAS 2.0 framework as an optional Legal Person Identification Data (LPID) attribute, meaning it can be included in EU digital identity and business wallet infrastructures.

For AI Agents, the LEI provides:

- Unambiguous, granular, identification of the organization behind an Agent,
- Global interoperability via an ISO standard Identifier and identity that is sector and jurisdiction agnostic,
- A neutral, open public infrastructure overseen by Regulatory Oversight Committee that includes the European Commission, ESMA, EBA, ECB and EIOPA.

This ensures that every AI Agent acting in commerce, payments, or infrastructure can be reliably linked to a real, verifiable organization.

vLEI: Cryptographically Verifiable Authority and Delegation

While the LEI identifies the organization, the verifiable LEI (vLEI) enables cryptographic proof of authority and delegation of granular rights and authorizations from organizations to digital credentials controlled by an Agent. The vLEI ecosystem provides cryptographically provable mandates, enabling an organization to digitally authorize

- a human
- a system
- or an AI Agent

to act on its behalf. This creates verifiable proof of who authorized the Agent and under what role or authority.

Machine-verifiable organizational roles: Using verifiable credentials, vLEI allows systems to verify roles:

- authorized employee
- procurement Agent
- payment authorization Agent
- supply chain automation Agent

These roles can naturally extend to AI Agents acting under defined mandates from the organization.

Trusted Agent-to-Agent interactions: When two AI Agents interact, each can verify:

- the organization behind the counterparty

- the mandate(s) of the Agent
- the validity of the credential chain

This enables secure Agent-to-Agent interactions in a multi-Agent coordination framework.

Annex B: PILOT CONCEPT 1: Supplier payments and treasury automation

This annex outlines a pilot concept (#1) for application of Agentic AI, developed during a multi-stakeholder workshop.

Context

Cross-border supplier payments remain operationally expensive, compliance-intensive and highly manual despite decades of payment automation. Agentic AI systems may be able to significantly improve arrangements for invoice validation, supplier onboarding, optimization of FX arrangements, settlement timing and liquidity management. The use case preserves a clear accountable-human chain and fits within existing regulatory frameworks.

Objective

Demonstrate cryptographically verifiable delegation of authority from a corporate treasury officer to an AI payment Agent capable of executing tightly constrained supplier-payment operations. The pilot would aim to assess whether (i) runtime policy enforcement can operate independently of credential issuance, (ii) delegated authority can remain interoperable across rails, (iii) auditability satisfies compliance requirements and (iv) vLEI can support machine-native authorization chains.

Figure 4: Use Case summary for payments automation

Use-Case Synopsis	
Core Participants	Buyer organization
	CFO or treasury officer
	AI payment Agent
	Supplier
	Payment Service Provider
	Policy engine operator
	Settlement rail
	GLEIF/vLEI ecosystem
Trust and Credential Chain	Legal Entity holds LEI and Legal Entity vLEI.
	CFO holds Official Organizational Role or ECR credential.
	CFO issues AMC/AARC credential to payment Agent.
	Agent presents delegated authority during payment execution.

Policy engine validates:	▪ mandate scope, ▪ sanctions, ▪ vendor allowlists, ▪ transaction limits, ▪ runtime conditions.

Operational Flow	In Scope	Out of Scope
1. Invoice received. 2. Agent validates invoice vs. PO 3. Delivery confirmation verified. 4. Agent composes payment intent. 5. Agent submits credential references. 6. Policy engine evaluates transaction. 7. Settlement rail executes payment. 8. Signed audit trail recorded.	✓ delegated authority, ✓ runtime policy enforcement, ✓ verifiable audit trail, ✓ cross-border B2B payment, ✓ rail-agnostic settlement	× autonomous corridor selection, × self-replicating Agents, × multi-Agent delegation chains, × autonomous treasury strategy.

We consider three issuer models for the attribute layer credentials, each with different operational and governance characteristics.

- **Model (a): Cloud providers as direct vLEI issuers.** AWS, Google Cloud, Microsoft Azure, and equivalent providers sign attestations of workload context (container ID, workload ID, image hash, region, instance metadata) directly as vLEI-credentialed issuers. This model is the architecturally cleanest because it brings the attestations closest to the source of truth. It requires cloud providers to enter the GLEIF vLEI governance framework as qualified issuers, which has both procedural and commercial dependencies that remain unresolved in the short term.

- **Model (b): Buyer organization as issuer, drawing on cloud provider metadata APIs.** The buyer organization operates the issuer infrastructure inside its own perimeter. The buyer queries cloud provider metadata exposed through standardized APIs (AWS Instance Metadata Service, Google Cloud metadata server, Azure Instance Metadata Service, Kubernetes Downward API) and signs the resulting attestations under its own vLEI. This model is the most realistic short-term option because it does not require cloud providers to enter GLEIF governance and uses APIs that already exist for operational purposes. It places the integrity responsibility on the buyer organization, which is a reasonable allocation of liability given the buyer is the party that benefits

from the Agent's actions. The deployment capabilities discussed in the Annex A supplement have made this model substantially more concrete than it was at the time of the workshop.

- **Model (c): Third-party attestation service as intermediary issuer.** A specialized service operates the issuer function on behalf of multiple buyer organizations, providing attestation infrastructure as a service. This model adds a trust hop relative to model (b) but creates a market for specialized issuers, which can professionalize the function and reduce per-buyer operational cost. The third-party introduces its own governance and liability questions that the paper should at least name.

The choice of issuer model has direct consequences for verification economics. Each additional issuer in the chain adds a Key Event Log and a Transaction Event Log that the policy engine must verify at runtime. Model (a) minimizes chain depth and verification cost. Model (b) adds one level. Model (c) adds two. For institutional B2B usage at moderate transaction volumes, all three are tractable. Annex D further discusses the verification cost implications across different use case volumes.

Settlement rail selection at the moment of payment could be partially independent of credential verification. The same operational flow can settle via SEPA, a stablecoin rail, a traditional wire transfer, or an internal account-to-account transfer, depending on the corridor, the amount, and the counterparty relationships. However, the reported consolidation of stablecoin infrastructure under major payment networks, (including the Stripe, Visa, and Mastercard collaboration reported on 3 June 2026 and the related Mastercard expansion of always-on stablecoin settlement capabilities the same week) would change the rail-agnostic positioning. Several major rails are now consolidating under the same infrastructure layer, which will reduce fragmentation but raise concentration questions.

Annex C: PILOT CONCEPT 2: Payments for data for due diligence

This annex outlines a pilot concept (#2) for application of Agentic AI that was developed during a recent multi-stakeholder workshop.

Context

AI Agents increasingly require real-time access to data from authentic sources as inputs to due diligence processes that can be cryptographically secured as audit trails for credit decisions. This can entail sourcing granular information including, credit data, sanctions screening, beneficial ownership information, trade records, registry data, and other commercial intelligence.

Current subscription-based models operated by data intermediaries or national platforms are poorly aligned with machine-native consumption of data in a controlled fashion that can drive payments in an economically viable manner at scale. Agentic AI managed processes could enable information from different sources to be sought out, accessed and verified and in parallel to generate nano-payments that remunerate authentic sources for data and veracity checks.

Objective

Demonstrate a vLEI-enabled ecosystem where AI Agents discover, authenticate and purchase signed data through machine-native payment protocols. The use case could be extended to support the DIFO/IFC invoice-finance workflows, autonomous due diligence, sanctions-screening Agents, procurement automation and tokenized trade-finance ecosystems

Figure 5: Use case synopsis for payments for due diligence data

Use-Case Synopsis	
Core Participants	**Data providers**
	Chambers of commerce
	Credit bureaus
	Registry operators
	AI buyer Agents
	Discovery marketplaces
	Curvy/X402 payment infrastructure
	GLEIF/vLEI ecosystem
Functional Model	1. Data provider exposes vLEI-verified API endpoint. 2. AI Agent discovers service. 3. Agent verifies seller identity. 4. Agent pays via X402-compatible infrastructure. 5. Seller releases signed data. 6. Provenance anchored into KERI event log. 7. Buyer receives verifiable data package.

Commercial Value	
For data providers:	For buyers:
• per-query monetization, • reduced unauthorized resale, • cryptographically provable provenance.	• machine-native access, • verified provenance, • reduced hallucination risk, • real-time enrichment. • Regulatory and Policy Relevance

The use case avoids many higher-risk concerns associated with fully autonomous consumer spending, speculative DeFi activities, or unrestricted AI financial operations. But it would test and demonstrate examples of machine-to-machine payments, identity verification, delegated authority, data provenance, and cross-border interoperability.

Annex D: The Economics of Verification Models

The economics of verification will influence design, adoption and usage. Verification is not a binary capability in which the system either supports verification or does not. In practice, verification has cost, and that cost scales with chain depth and transaction volume. Different use cases sit at very different points on the cost curve. This is worth making explicit, because architectural choices that work for institutional B2B may not work for high-frequency machine-to-machine micropayments, and vice versa.

The architectural choice between attribute layer mechanisms (see Section 6) has real operational consequences that vary with transaction volume. Edge-chained credentials, schema variants, and external attribute logs each impose different verification costs per transaction. The same is true for the issuer model choice discussed in Section 6 and the Annex A supplement. Each additional issuer in the credential chain adds verification work. The economics of these choices will only become visible once the architecture is placed against representative transaction volumes.

Two scenarios discussed in the April workshop make the trade-offs visible at opposite ends of the spectrum.

(i) Agent-mediated Uber rides, grocery, or consumer-card-rail transactions.

These transaction examples were raised by Professor Hardjono as low-risk consumer scenarios. Their order of magnitude in the United States alone is approximately 28 million such transactions per day. At a chain depth of around four credentials per transaction (legal entity, role credential, Agent mandate, transaction authorization), full ACDC verification cost is dominated not by signature checking but by how often issuer key and revocation state must be resolved from witnesses.

Verification speed and computing power matter. With key state warm in cache, a depth-four chain verifies in roughly a millisecond, and the consumer-scale footprint is under one server-day of verification compute per day of operation. With every chain resolved cold from witnesses, on the order of several seconds per chain per the workshop estimate, the same volume rises into the thousands of server-days. Realistic deployments sit between these bounds, in the order of tens to low hundreds of server-days, set by cache-hit rate and the acceptable staleness of revocation data. The precise operating point should be benchmarked. The architectural point is that consumer-scale verification is tractable only under a caching strategy, which in turn bounds how fresh revocation can be. This is the use case the workshop discussion characterized as the natural starting point for consumer-facing Agentic payments, on the grounds that verification economics are tractable and existing card-rail liability frameworks apply.

(ii) Agent-to-Agent micro-payment

Curvy outlined a scenario based on the DIFO[22] protocol developments led by Finspot and IFC. In that scenario, data is collected and verified by multiple sources to support the due diligence process for invoice finance and factoring. This is currently managed via legacy data controls and integrations but could be managed in future by AI Agents to pay for data feeds, API access, compute time, and other machine-consumed resources in real time. This might be better managed and more economical if managed without direct human supervision of each transaction. The relevant order of magnitude is qualitatively different. If one million autonomous research Agents operate globally and each performs ten thousand micropayments per day, daily transaction volume reaches approximately 350 times the consumer-rail baseline. At four credentials per transaction and current verification cost assumptions, the per-transaction verification cost approaches or exceeds the transaction value itself.

Recent data from the x402 protocol on the Base blockchain provides an empirical anchor for the Agent-to-Agent end of the spectrum. Base reported approximately 3.1 million x402 transactions in the 30 days ending 29 May 2026, with approximately $1.2 million in total value transferred over the same period.[23] Sellers grew by 23 percent and buyers by 37 percent month over month, indicating a high growth rate from a baseline that is already non-trivial. The current 30-day volume corresponds to approximately 103,000 transactions per day across a single rail and a single chain, which is well below the speculative one-million-Agent baseline but is already meaningful and is growing rapidly. The average transaction value is approximately $0.39, which makes verification cost economics critically sensitive to chain depth and to the cost per verification step.

Implications for design and economics:
The position that follows from these scenarios. The vLEI and ACDC ecosystem is ready for institutional B2B use cases and for consumer-card-rail-equivalent scenarios at current verification cost assumptions. Edge-chained credentials work well at these volumes. The open architectural question that becomes more pressing as Agentic micropayment volumes grow is how to make verification economics work at machine-cadence transaction volumes. This is not a limitation of what vLEI and ACDC deliver today. It is a known frontier that requires further architectural work, including the assessment of external attribute log

[22] Digital Invoice Finance Origination. This initiative provides an open protocol reference framework for tokenization and digital verification of individual data points and their provenance from authentic sources in order to provide for fully tokenized, auditable due diligence trails of financial assets.
[23] The public x402.org dashboard and Chainalysis report a substantially higher figure for the same protocol over a comparable period, on the order of 72 to 75 million transactions and approximately $24 million in value. The 3.1 million figure used here is the narrower Agent-payment count reported by Base.

mechanisms, batched verification optimizations, and the role of caching and pre-verification at the policy engine level.

Implications for the architectural choices outlined in **Section 6**

Of the three attribute layer options discussed, the relative trade-offs map onto verification economics as follows.

- **Edge-chained credentials** maintain native verifiability but incur chain traversal cost that grows linearly with chain depth. For institutional B2B at the consumer-rail scale, this cost is acceptable. But for machine-cadence micropayments, it dominates the transaction.

- **Schema variants** reduce verification cost to a single credential check but at the cost of schema rigidity, which makes them suitable for use cases with a well-defined attribute taxonomy.

- **External attribute logs** reduce verification cost to a single hash check at the cost of trust in the external infrastructure, which makes them potentially suitable for high-frequency scenarios where verification cost dominates other considerations.

Of the three issuer models described in Section 6 and the Annex A supplement, each adds verification overhead in proportion to chain depth.

- Model (a) with cloud providers as direct vLEI issuers minimizes chain depth.

- Model (b) with the buyer organization as additional issuer adds one level.

- Model (c) with a third-party attestation service adds two levels.

For institutional B2B at moderate transaction volumes, all three models are tractable. For high-frequency machine-to-machine scenarios, the choice has economic weight.

The policy-engine operator bears the verification cost. The operator question raised in Section 6 has a direct economic dimension that belongs here, because whoever runs the runtime policy engine is also the party that absorbs the per-transaction verification cost set out above. A bank operating the policy engine internalizes chain resolution and revocation checks as part of its existing compliance processing, where the marginal cost is small relative to the transaction. A settlement rail operating the engine must amortize that cost across high transaction volumes at low per-transaction value, which is precisely the machine-cadence regime where verification economics are most sensitive. An identity-layer operator sits between these positions. The choice of operator therefore interacts with the

caching strategy: the operator best placed to maintain warm key and revocation state at the relevant volume is also the operator for whom the economics are most favorable.

Deployment profile moves the operating point on the cost curve. The split between warm-cache and cold-from-witness verification described above is, in large part, a deployment-profile question. A self-hosted verifier operating within an institutional perimeter can maintain warm key state, pre-resolve revocation for known counterparties, and control its own caching policy, which places it toward the favorable end of the cost range for institutional B2B volumes. A cloud-hosted verifier serving many tenants may achieve good cache-hit rates through scale but introduces shared-infrastructure trust and staleness considerations that bear on how fresh revocation data can be. The deployment choice discussed in Section 6 thus has an economic counterpart: it does not change the shape of the cost curve, but it determines where a given deployment realistically sits on it.

References

Aldasaro, Iñaki and Desai, Ajit. AI Agents for cash management in payment systems. BIS Working Papers No 1310, November 2025.

Australian Productivity Commission. Harnessing data and digital technology. Inquiry Report No. 111, December 2025.

Birch, David and Hoffart, Jelena. Know your Agent: Enabling autonomous financial services. Journal of Digital Banking Vol. 10, No. 2, 2025.

Compliance-Aware Agentic Payments on Stablecoin Rails. arXiv, 29 April 2026.

Deutsche Bank and Clifford Chance. Convergence of AI and DLT. 2025.

Digital Regulation Cooperation Forum (of the United Kingdom). The Future of Agentic AI, Foresight Paper. 31 March 2026.

Elliptic Curve Pairing Stealth Address Protocols. arXiv:2312.12131; and Post-Quantum Stealth Address Protocol. eprint.iacr.org/2025/112.

Fenwick. Is 2026 the Year of Agentic Payments? 21 April 2026.

FIDO Alliance. FIDO Alliance to Develop Standards for Trusted AI Agent Interactions. 27 April 2026.

Google. Announcing Agents-to-Payments (AP2) protocol, 2025. As cited in the discussion paper literature review.

Gosmar, Diego; Attwater, David; Coin, Emmett; Dahl, Deborah. AI Multi-Agent Interoperability Extension for Managing

IETF draft on AI Agent Authentication and Authorization, as summarized in the discussion paper literature review

Davidovic, Sonja, and Hervé Tourpe. IMF. How Agentic AI Will Reshape Payments. 24 April 2026.

Infocomm Media Development Authority of Singapore (IMDA). Model AI Governance Framework for Agentic AI. Version 1.0, 22 January 2026.

Reed, Chris. Autonomy, Responsibility and Agentic AI. As discussed in the literature review summary.

South, Tobin; Marro, Samuele; Hardjono, Thomas et al. Authenticated Delegation and Authorized AI Agents. arXiv:2501.09674v1, 16 January 2025.

Terminal 3 (2026). Agent Command: Decoupling Reasoning and Execution for Secure Agentic AI. Terminal 3 Pte. Ltd. Available at: terminal3.io

X402 protocol. Coinbase / Cloud Native Computing Foundation. HTTP 402 Payment Required: machine-native payment standard for API monetization.