Beyond Single-Enterprise ZTA: Multi-Trust Architectures for Authorised Agentic Actors in Open, Cross-Domain Ecosystems

GSMA Greater China's M-Trust Proposal and Its Alignment with Business Wallets, W3C Verifiable Credentials, Decentralized Identifiers, Post-Quantum Cryptography and Distributed Ledger Technology for Trusted Agentic Networks

An architectural and geopolitical analysis of the GSMA Greater China draft

Dr. Carsten Stöcker
Spherity GmbH

Research and evidence cut-off: 16 September 2026

Agentic networks cannot operate reliably, safely and securely on the basis of unverified, naïve confidence in an agent. They require verifiable evidence of the legal person behind the agent, the mandate under which it acts, the AI service and runtime state involved, the evidence supporting the action, and the policy decision that authorised execution.

Abstract

This article applies a qualitative architecture and standards gap analysis to the GSMA Greater China paper *Trust Paradigm Evolution for Agentic Networks*. The paper presents M-Trust as an extension of Zero Trust for autonomous, multi-party and cross-domain environments [1]. Its three dimensions—Multi-party trust, Agentic Intent-aware trust and Cross-domain trust—provide a useful systems model for future sixth-generation (6G) networks and artificial-intelligence-native infrastructures. The draft explicitly identifies World Wide Web Consortium (W3C) Decentralized Identifiers, Distributed Ledger Technology and Post-Quantum Cryptography (PQC). It describes lattice-based cryptography, lightweight implementations and a hybrid cryptographic architecture. It cites no NIST FIPS 203, FIPS 204 or FIPS 205 and no named Chinese PQC standard. W3C Verifiable Credentials, Business Wallets and legally grounded delegation remain unspecified.

The geopolitical analysis interprets the publication as early agenda setting. Bruegel characterises China's AI strategy as combining domestic control and technological autonomy with global diffusion, open models and international standard setting [2]. The separate GSMA *Mobile AI* paper explicitly calls for globally unified technical standards developed through international standards organisations and industry alliances [3]. M-Trust fits this observable pattern: a conceptual vocabulary developed by China-based public, industrial and academic actors enters a global forum while United States and European agentic-authority profiles remain under development.

A comparison with Anthropic's *Zero Trust for AI Agents*, the WE BUILD non-paper, European data-space initiatives such as Manufacturing-X, Catena-X and energy data-X identifies a converging control pattern [4–9]. Anthropic promotes cryptographically rooted

agent identities, task-scoped permissions, protected memory and agent-oriented defensive controls at a more abstract enterprise-security level. WE BUILD connects European Digital Identity and Business Wallets with verifiable credentials, provable mandates, mutual authentication and chains of proof. European data spaces use a common trust infrastructure and legal-person identity across domain-specific governance frameworks; Manufacturing-X applies this logic to interoperable industrial value chains. These publications and initiatives are organisational frameworks rather than one harmonised standard. Their common pattern combines verifiable identity, bounded authority, contextual policy, external enforcement and reconstructable evidence.

The M-Trust scenarios primarily model a technical relationship in which an agent acts on behalf of a user and is evaluated through user, device, network and behavioural evidence. They do not specify legal-person identity, corporate representation or organisational accountability. This article therefore distinguishes a B2C branch, in which a natural person authorises a shopping or payment agent, from a B2B/B2G branch, in which an agent is bound to an identified legal person through a verifiable representation and delegation chain. The legal-person extension is being researched, tested and evaluated in Europe through European Digital Identity and Business Wallet work, WE BUILD, Manufacturing-X, Catena-X and related Spherity research [6–9]. It targets cross-company industrial actions whose economic, compliance and cyber-resilience consequences require accountable organisational authority.

This convergence makes M-Trust useful beyond telecoms. Industrial data spaces, Digital Product Passports (DPPs), supply-chain programmes and Manufacturing-X require organisations and agents to traverse several trust domains while preserving issuer provenance and local verifier sovereignty. A European Business Wallet can present credentials from company-register, market-role, conformity-assessment, regulatory and sector-governance domains. The verifier composes those credentials into transient cross-domain trust for a particular transaction without merging the source domains or treating upstream trust as universally transitive.

This article contributes a target control plane for economically or legally material agent actions. A Business Wallet binds an agent to an identified legal person and governs delegation. W3C Verifiable Credentials carry authoritative identity, role, licence, mandate and AI-service evidence. DIDs support identifier resolution and cryptographic control. Verifier policy decides whether the evidence authorises a requested action. PQC migration protects the public-key trust fabric. A governed verifiable data registry (VDR)—implemented through a conventional registry, Distributed Ledger Technology, the European Blockchain Services Infrastructure (EBSI) or a qualified electronic ledger service—supports shared status and integrity functions according to risk and legal qualification. For IPCEI-AI candidate projects, the same pattern can connect the control plane to a data-plane verifiable evidence graph: a machine-resolvable graph of signed provenance, conformity, testing, runtime and action evidence. The resulting architecture aligns M-Trust with Spherity's Trusted AI model: verifiable authority + verifiable evidence + policy enforcement + bounded autonomy.

Keywords. Agentic networks; M-Trust; Zero Trust; B2C agentic commerce; verifiable intent; legal-person identity; Business Wallets; W3C Verifiable Credentials; Decentralized Identifiers; verifiable authorisation; Trusted AI; Post-Quantum Cryptography; Distributed Ledger Technology; EBSI; Manufacturing-X; vLEI; data spaces

Contents

Key Findings **5**

1 Introduction **6**

1.1 Research Questions 7

1.2 Method and Source Corpus 7

1.3 Scope Limitations and Author Positionality 8

1.4 Evidence Classes and Normative Language 8

2 Conceptual Framework: M-Trust, Zero Trust and Domain Instantiations **8**

2.1 Telecommunications Examples for M-Trust 10

2.2 Human–Agent Authorisation and Verifiable Intent in B2C Agentic Commerce . . 10

2.3 Energy Data-X and Cross-Sector Trust Domains 12

3 Comparative Results across China, Europe and the United States **14**

4 Standards Requirements and Verifiable Trust Chains **16**

4.1 Documented Standards Coverage 16

4.2 Standards Comparison: M-Trust and Manufacturing-X 17

4.3 Functions in a Verifiable Trust Chain 19

5 Organisational Identity and Verifiable Authority **19**

5.1 Business Wallets as the Organisational Control Plane 19

5.2 Decentralized Identifiers and Identifier Governance 20

5.3 Verifiable Credentials as Semantic Trust Objects 21

5.4 Verifiable Authorisation Chains and Assurance Graphs 23

6 Cryptographic and Registry Infrastructure **24**

6.1 Post-Quantum Cryptography and Cryptographic Agility 24

6.2 Distributed Ledger Technology and Bounded Registries 28

7 Applications in Data Spaces, Digital Product Passports and Trusted AI **29**

7.1 Data-Space Identity and Protocol Layers 29

7.2 Manufacturing-X and the Catena-X Trusted AI Path 30

7.3 IPCEI-AI Outlook: From Control Plane to Evidence Data Plane 31

8 Alignment with the Spherity Trusted AI Architecture **32**

8.1 Work with Bundesanzeiger Verlag and the German Company Register 32

8.2 Alignment with GLEIF for Cross-Jurisdiction vLEI Trust 33

9 Reference Architecture for Verifiable Agentic Networks **33**
9.1 Multi-Trust-Domain Governance and Transient Trust 34
9.2 End-to-End Authorisation Transaction 35

10 Discussion: Geopolitics, Standardisation and Industrial Adoption **35**
10.1 Source Composition and Analytical Boundary 35
10.2 Vocabulary Power 36
10.3 Distribution Power 36
10.4 Infrastructure Power and China Standards 2035 36
10.5 A Multipolar Interoperability Response 36
10.6 Strategic Consequences of Missing Trust Infrastructure 37

11 Research and Implementation Priorities **37**
11.1 Semantic and Governance Profiles 38
11.2 System Integration and Cryptographic Resilience 38
11.3 Validation and Empirical Research 38

12 Conclusion **39**

Declarations **40**

References **40**

Key Findings

1. **W3C DIDs appear explicitly while W3C Verifiable Credentials remain unspecified.** M-Trust names the W3C DID standard and describes standardised trust credentials, identity credentials, on-chain credentials and claims. It neither names the W3C VC Data Model 2.0 nor defines issuer–holder–verifier roles, presentation semantics, holder binding, status or privacy profiles [1, 10, 11].

2. **M-Trust is user-centred rather than legal-person-centred.** Its worked identity pattern binds a dedicated agent identity to a user and evaluates whether the agent and its behaviour are authorised. The 'user' remains a technical identity rather than a legally established natural-person identity. Legal-person identity, corporate representation, Business Wallet governance and organisational liability remain outside the specified architecture. European work extends this pattern towards accountable B2B/B2G authority through Business Wallets, organisational credentials and verifiable delegation [1, 6–9].

3. **M-Trust extends Zero Trust across three material gaps.** Multi-party evaluation, intent-aware authorisation and cross-domain trust address the limits of resource-centric enterprise architectures when autonomous agents coordinate across networks and organisations.

4. **Trust assessment and legal authorisation serve different decision functions.** Reputation, behavioural scoring, consensus and third-party endorsements can inform risk. Legal-person identity, governed roles and a valid delegation establish authority. A receiving domain remains the sovereign decision point.

5. **Business Wallets and verifiable authorisation chains fill the principal institutional gap.** A Business Wallet connects legal-person identity, representatives, agents, approval rules, credentials, revocation, recovery and audit. A material action requires a current and scoped chain from an authoritative source to the legal person, agent, task, policy decision and action receipt.

6. **The GSMA papers provide PQC direction without a standards profile.** M-Trust names PQC and lattice-based cryptography and anticipates hybrid cryptography. Its references cite NIST SP 800-207 for Zero Trust while citing no NIST FIPS 203, FIPS 204 or FIPS 205 and no named Chinese PQC standard. The Chinese draft mentions the SM series in a separate classical software-root discussion. The Mobile AI paper contains no PQC passage [1, 3].

7. **PQC migration concerns the complete public-key trust fabric.** Interoperability requires coordinated profiles across issuers, wallets, verifiers, trust lists, status services, DID methods, verifiable data registries, secure channels, timestamps and archived evidence.

8. **DLT has a bounded and optional role.** Ledgers can support shared ordering, integrity anchors, identifier operations, credential status and audit commitments. Issuer governance establishes claim authority; verifier policy establishes fitness for purpose.

9. **The comparative and geopolitical findings concern agenda-setting proposals.** M-Trust, Anthropic, WE BUILD and Manufacturing-X independently share patterns of strong identity, bounded permissions, contextual policy and accountable execution. Publication of M-Trust in English through GSMA gives a China-developed vocabulary direct visibility among global telecom decision-makers. The documents provide evidence of architectural convergence, rather than evidence of coordinated regional policy [1–3, 5–9].

10. **Manufacturing-X and sector data spaces provide an industrial implementation path.** Their federated trust models can combine a Business Wallet with credentials from company-register, market-role, conformity-assessment and regulatory domains. The result is transaction-specific cross-domain trust across energy, manufacturing, mobility, DPP and supply-chain use cases [7–9, 12].

11. **Regional PQC trajectories remain distinct.** The United States has finalised NIST standards and federal migration profiles; the European Union coordinates a risk-based transition while Germany maintains additional implementation and sovereignty research; China operates an indigenous next-generation commercial-cryptography selection programme. M-Trust selects none of these standards paths [1, 13–19]. Diverging PQC roadmaps are an unresolved interoperability issue for global trust architectures.

12. **EBSI and qualified electronic ledgers require separate treatment.** EBSI provides a European shared trust-layer and VDR pattern. Qualified electronic ledger is an eIDAS legal status tied to a qualified trust service provider and statutory requirements. A locally operated node can reduce live dependency on a single registry intermediary while preserving the need for issuer, governance, software and consensus assurance [20, 21].

13. **Trusted AI adoption depends on production-grade trust infrastructure.** Regulated industry and critical infrastructure require machine-verifiable legal authority, compliance evidence, functional-safety assurance and cyber-resilience. Delay in this control infrastructure creates adoption latency, ecosystem fragmentation and a strategic late-mover disadvantage.

1 Introduction

Agentic systems can select tools, exchange data, initiate transactions and coordinate with other agents across organisational boundaries. These capabilities expand the scope of security decisions. A relying party must evaluate the technical identity of the agent, the natural or legal person that bears responsibility, the mandate for the proposed action, the current state of the AI service and the evidence available at decision time.

Two application classes require separate treatment. In B2C agentic commerce, a natural person authorises an agent to search, select or transact within stated constraints. The core authority question concerns the person's authenticated instruction, its scope and the evidence of assent to the resulting transaction. In B2B and B2G settings, the agent acts for a legal person. The verifier must therefore establish the organisation, the representative or system that delegated authority, the applicable internal approval rules and the agent's task-bounded mandate. Spherity's joint research focuses on this B2B/B2G branch because cross-company industrial transactions, regulated data exchange and critical-infrastructure operations concentrate substantial macroeconomic value and systemic risk. Reusable legal-person trust infrastructure can therefore support industrial competitiveness and cyber-resilience at the same time.

The M-Trust draft provides a systems model for these decisions in mobile and AI-native networks. Its Multi-party, Agentic Intent-aware and Cross-domain dimensions extend Zero Trust towards autonomous interaction [1]. This article examines the standards and governance required to translate that model into verifiable cross-organisational authority. It contributes a standards gap analysis, a normative authorisation model and a target architecture that connects M-Trust with Business Wallets, data spaces and Spherity's Trusted AI research.

The argument proceeds in four stages. Sections 2 and 3 establish the analytical framework and compare the source architectures. Sections 4 to 6 identify the missing identity, authority,

cryptographic and registry components. Sections 7 to 9 apply those components to data spaces, Digital Product Passports and Trusted AI and then integrate them into a reference architecture. Sections 10 to 12 examine geopolitical implications, implementation priorities and the resulting conclusions.

1.1 Research Questions

This study focuses on B2B agentic networks, including relevant B2G interactions, in which enterprise agents act across organisational, sectoral and jurisdictional boundaries. B2C agentic commerce provides broader context for dual human–agent identity, verifiable intent and emerging agentic payment models.

The analysis addresses five questions:

1. Which identity, credential, wallet, cryptographic and registry components appear in M-Trust?
2. Which additional components support legally attributable enterprise-agent actions across organisational boundaries?
3. How can M-Trust align with Business Wallets, data spaces and Spherity's Trusted AI architecture?
4. Which governance choices follow for Europe, North America and Asia-Pacific?
5. Which architecture patterns recur across the China-origin M-Trust proposal, Anthropic's United States corporate framework, the European WE BUILD non-paper and Manufacturing-X implementation ecosystems?

1.2 Method and Source Corpus

The method is a qualitative architecture and standards gap analysis. The primary corpus comprises the English and Chinese M-Trust drafts, the GSMA *Mobile AI* paper and the Spherity *Trusted AI* architecture deck. The comparison uses NIST Zero Trust Architecture, Anthropic's *Zero Trust for AI Agents*, the WE BUILD non-paper, W3C DID Core, W3C Verifiable Credentials Data Model 2.0, the Data Spaces Support Centre Blueprint, Manufacturing-X and Catena-X materials, NIST post-quantum standards, European and German PQC sources, Chinese commercial-cryptography programme materials, European Business Wallet and GLEIF publications, EBSI and eIDAS ledger provisions, and related Spherity Research.

The method evaluates the two branches separately. The B2C analysis compares first-party public materials from Google, Mastercard, Visa, Amazon and Apple for human authentication, agent recognition, delegated instructions, transaction confirmation and evidence portability. It then applies the UNCITRAL Model Law on Automated Contracting as a reference point for the distinction between technical proof and legal effect. The B2B/B2G analysis examines legal-person identity, representation, organisational policy and cross-domain credentials in European wallet, data-space and registry initiatives. This separation avoids treating a consumer account, an organisational identity and a legally effective mandate as equivalent evidence.

The review tests whether each architectural claim is stated by a source, follows as an inference or originates in this article's normative proposal. The English and Chinese M-Trust drafts were compared for the DID, credential, DLT, classical-cryptography and PQC passages. The *Mobile AI* paper was searched separately for PQC and standardisation language. Standards claims were

checked against the relevant primary specifications. The analysis evaluates documentary content and conceptual consistency. It provides no empirical evidence of performance, deployment security, legal qualification or political coordination.

1.3 Scope Limitations and Author Positionality

The source documents differ in status. M-Trust is a draft reference architecture and expressly disclaims status as a mandatory specification [1, p. 29]. Anthropic provides corporate security guidance [5]. WE BUILD provides a consortium non-paper [6]. These sources therefore support a comparison of published architecture patterns. They do not establish national or regional consensus.

The author is the founder and chief executive of Spherity and a named co-author of the WE BUILD non-paper. Spherity publications and the internal architecture deck are treated as design inputs from the joint research programme, rather than independent validation of that programme. Geopolitical claims are limited to source composition, publication channel, documented strategy and the plausible effects of agenda setting.

1.4 Evidence Classes and Normative Language

The analysis separates three evidence classes:

- **Documented design:** an element stated in the source paper.
- **Architectural inference:** an implication derived from the interaction of stated components.
- **Normative proposal:** a requirement for the target architecture developed in this article.

The uppercase terms **MUST**, **MUST NOT**, **SHOULD** and **MAY** follow BCP 14 [22,23]. They identify normative proposals from the joint research presented here. GSMA-derived directions appear as descriptive source statements and carry a citation to [1] or [3]. Each normative block also states its source explicitly.

2 Conceptual Framework: M-Trust, Zero Trust and Domain Instantiations

NIST Zero Trust Architecture moves security away from static network perimeters and centres protection on users, assets and resources. It treats subject and device authentication and authorisation as discrete functions performed before a session to an enterprise resource is established [24]. M-Trust retains identity-centric security, least privilege, continuous verification and dynamic risk control, then adds three dimensions for agentic networks [1].

Figure 1 maps this three-stage progression and locates the M-Trust decision loop within the proposed verifiable-authority layer. The figure depicts the trust mechanisms for an authorised enterprise agent.

Figure 1 maps this three-stage architectural progression and situates the M-Trust decision loop within the proposed verifiable-authority layer. It depicts the trust mechanisms for an authorised enterprise agent acting on behalf of a legal person. In a B2C scenario involving a human user's agent, the agent operator could provide verifiable agent-assurance claims through its

Business Wallet, while the human user could use a natural-person wallet to provide task-specific authorisation and evidence of intent. Together, these evidence paths would support M-Trust's dual human–agent identity model. The legal effect of the user's declaration remains governed by applicable law. This B2C variant lies outside the scope of the figure.

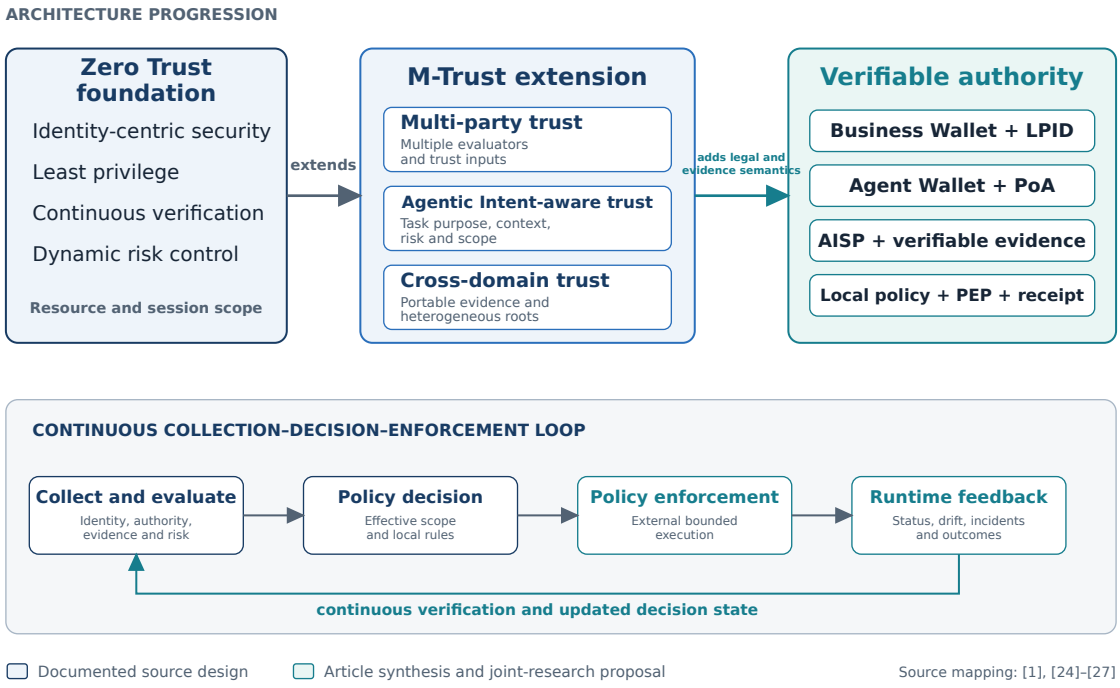

Figure 1: Progression from Zero Trust through M-Trust to verifiable authority for B2B use cases. Source: Author's synthesis based on [1, 24–27].

M-Trust dimension	Architectural contribution	Trusted AI interpretation
Multi-party trust	Multiple evaluators, trust models, endorsements and consensus inputs	Multi-issuer evidence, trust registries, assurance graphs and counterparty risk signals
Agentic Intent-aware trust	Task objective, execution context, resource request and risk enter the authorisation decision	A bounded mandate, declared purpose, tool scope, value limit, runtime evidence and local policy enter each material decision
Cross-domain trust	Heterogeneous trust roots, portable credentials and dynamic trust-chain maintenance	Interoperable credentials and presentations, issuer governance, status, semantic profiles and verifier-sovereign policy

M-Trust also defines a useful closed loop [1, pp. 11–17]:

trust collection and evaluation → policy decision → policy enforcement → runtime feedback

This loop matches modern policy decision point and policy enforcement point patterns. It also creates a place for continuous AI assurance. Behavioural drift, prompt-injection indicators, tool-call anomalies, credential status changes and new threat intelligence can alter the decision state during a task.

M-Trust therefore extends the *operational scope* of Zero Trust. The Spherity architecture proposed here adds the *legal and evidentiary semantics* required when an agent enters a contract, accesses regulated data, initiates a payment, changes an industrial process or controls a physical system. This distinction supplies the logical bridge from risk-aware access control to verifiable authority.

2.1 Telecommunications Examples for M-Trust

The GSMA draft grounds M-Trust in concrete telecom scenarios. These examples remain conceptual use cases rather than reported production deployments [1, pp. 23–28].

GSMA scenario	**M-Trust problem**	**Verifiable-authority interpretation**
Agent requests cross-domain resources for a user	The network must verify both the agent–user binding and the requested behaviour	A human or legal-person identity credential, agent identity and task-scoped mandate establish accountable authority
Access, security, demand and scheduling agents form temporary teams	Agents from separate domains may lack a pre-established relationship and disband when the task ends	Portable credentials, short-lived presentations and local policy create transient trust for the collaboration
Agent establishes a cross-operator high-definition video call	The task invokes access scheduling, bandwidth, encryption and monitoring across domains	A task–permission–invocation binding narrows rights as the workflow advances and produces signed receipts
In-vehicle, mobile-terminal and cross-domain service agents roam	The technical trust anchor and identity continuity must survive domain changes	SIM/eSIM, hardware roots and software roots establish technical provenance; organisational credentials establish legal authority
Operators share spectrum or collaboration state	Several parties require common ordering, integrity and auditability	A governed ledger can record minimal shared state and commitments while credentials and sensitive data remain off ledger

M-Trust's proposed dual human–agent identity[1], dynamic allocation and post-task reclamation of permissions closely resembles a short-lived agent Power of Attorney enforced at each invocation. Its SIM/eSIM model offers an operator-endorsed technical root for mobility. The legal person, organisational role and right to commit resources still require an institutional identity and delegation layer. The same separation applies to the draft's hardware-based and software-based roots of trust: they can attest a device or runtime while a governed credential chain establishes the authority represented by that runtime.

2.2 Human–Agent Authorisation and Verifiable Intent in B2C Agentic Commerce

M-Trust's dual user–agent identity is most directly applicable to B2C agentic commerce. The principal is a natural person who delegates search, selection or purchase activity to an agent. The

[1]Human–agent identity pattern can be implemented through delegation or authorisation credentials whose creation or presentation is approved through a natural-person wallet, such as an EUDI Wallet. Related representation and delegation mechanisms are being researched in WE BUILD and German SPRIND-supported EUDI Wallet projects. Payment-sector initiatives are also examining verifiable intent and agent-initiated payments for human-controlled shopping agents. A common credential and policy profile across these domains remains a research and standardisation priority.

relying merchant or payment provider must distinguish an authorised agent from an anonymous bot, determine which user stands behind the interaction and verify the scope of the user's instruction. By contrast, the B2B/B2G branch examined in this article requires a legal-person identity and a representation chain in addition to a technical user–agent binding.

Public initiatives from Google, Mastercard, Visa, Amazon and Apple address different layers of this problem [28–33]. The term 'verifiable intent' is used here as an analytical category for evidence that links a natural person's instruction to an agent and a resulting action. It does not imply that all five providers implement the same protocol, evidence model or degree of portability.

Provider and initiative	**Primary mechanism**	**Authorisation evidence**	**Analytical boundary**
Google: Agent Payments Protocol (AP2) [28]	Open, payment-agnostic protocol using an Intent Mandate and a Cart Mandate	Cryptographically signed instructions, constraints and cart approval; the published model uses verifiable credentials and supports human-present and delegated purchases	Protocol evidence requires applicable rules for identity assurance, liability, consumer rights and legal effect
Mastercard: Verifiable Intent and Agent Pay [29]	Protocol-agnostic record that links identity, intent and action; planned integration with Agent Pay intent APIs	Tamper-resistant evidence of cardholder authorisation, specific instructions and the agent–merchant interaction; selective disclosure and future VC integration are described	The cited publication describes an open specification and prospective integrations rather than a uniform legal-authorisation regime
Visa: Trusted Agent Protocol [30]	Signed requests from approved agents, bound to a merchant domain and operation	Agent identity, user permission, freshness, session context and transaction-specific authorisation; optional consumer and payment identifiers	The public material emphasises trusted-agent recognition and action binding; it reports that the product remains in development and deployment
Amazon: Buy for Me [31]	Platform-mediated purchase from an external brand site after an Amazon checkout confirmation	User confirmation of order details, address, fees and payment method; encrypted payment and shipping data are supplied to the merchant	The evidence remains within Amazon's account and workflow rather than a published, portable cross-platform mandate
Apple: App Intents and StoreKit [32, 33]	Structured app actions exposed to system services, combined with platform-secured in-app purchase processing	Typed action parameters, optional confirmation interactions and StoreKit transaction records within the Apple ecosystem	The cited materials define application actions and purchase processing; they do not define a portable cross-merchant proof of delegated purchase intent

Google AP2 and Mastercard Verifiable Intent come closest to a portable mandate model: both preserve a cryptographic record of the user's instruction, with AP2 separating an initial intent from the approved cart and Mastercard linking identity, intent and action [28, 29]. Visa binds a recognised agent and its authorisation to the merchant domain and requested operation [30]. Amazon uses explicit confirmation inside a closed account and checkout workflow [31]. Apple provides a structured action framework and secure purchase layer; in the cited materials, the term 'App Intent' denotes an application capability rather than a portable expression of legal intent [32, 33]. Where the action becomes an input to a larger workflow, its signed receipt can

form a node in the verifiable evidence graph developed later in this article [34].

Technical verifiability and legal intent require separate tests. A signature, credential or platform record can establish integrity, authentication, scope, timing and the content of an instruction. Applicable law determines whether those facts amount to legally effective assent or payment authorisation and how capacity, attribution, pre-contractual information, mistake, unexpected outcomes, withdrawal, unfair terms, liability and remedies apply. The UNCITRAL Model Law on Automated Contracting provides a reference framework for recognising automated contracting, attributing system outputs and addressing unexpected outcomes. As a model law, it requires implementation in applicable law and does not itself harmonise consumer-protection, payment or agency rules [35].

The resulting B2C chain is natural person → authenticated wallet or account → scoped intent or delegation → agent → cart or payment action → receipt. It operationalises the human–agent side of M-Trust. The B2B/B2G chain adds legal-person identity, representative authority, organisational approval policy and cross-domain credentials. A consumer confirmation and a corporate authorisation chain therefore share a control pattern while serving different legal and governance functions.

Normative source — joint research: A B2C verifier SHOULD evaluate the natural person's authentication, the agent identity, the scope and freshness of the instruction, transaction-specific confirmation, status and the resulting receipt. A B2B/B2G verifier SHOULD additionally establish the legal person, representative authority, organisational approval rules and the complete delegation chain. These requirements extend the M-Trust pattern; they are not stated GSMA requirements.

2.3 Energy Data-X and Cross-Sector Trust Domains

The energy data-X research provides a sectoral implementation analogue. Its reported case connects organisational and market-role evidence with a Business Wallet and connector-side access control for federated energy-data exchange [12]. The next architectural phase applies the same trust pattern to AI agents that discover data, negotiate usage conditions, generate forecasts, coordinate battery energy-storage systems and initiate bounded market or operational actions.

The design separates the following authoritative trust domains. The credential forms below are target profiles from the joint research; the underlying registers, codes and certifications are existing authoritative sources.

Trust domain and authoritative source	**Existing authority**	**Proposed machine-verifiable evidence**	**Example decision**
German Company Register, operated by Bundesanzeiger Verlag	Official register information concerning the legal person and representation	Legal Person Identity and representation credentials	Is this agent bound to the registered company and an authorised delegator?

Trust domain and authoritative source	Existing authority	Proposed machine-verifiable evidence	Example decision
BDEW Codes and Services (Energie Codes und Services GmbH)	This wholly owned BDEW subsidiary allocates codes that uniquely identify a participant and its role in the German electricity market; it also operates validation interfaces and acts as the German electricity-market issuing office for Energy Identification Codes	Authorised market-participant and market-role credentials derived from validated code records	May this organisation act as the claimed supplier, grid or market participant?
Testing, Inspection and Certification (TIC) domain	Accredited conformity-assessment bodies (CABs), including TÜV organisations, assess and certify management systems and products	Certificate credentials, including ISO/IEC 42001 AI-management-system certification where the certification scope applies	Does the relevant organisation or AI service carry current conformity evidence from an accepted body?
Market Surveillance Authority (MSA) domain	German and European public authorities exercise legally assigned market-surveillance functions	MSA role and jurisdiction credentials issued under the applicable governmental trust framework	May the requesting official inspect restricted product, DPP or compliance evidence?
German Energy Agency (dena) expert domain	dena operates the federal Energy Efficiency Expert List for funding programmes	Authorised energy-adviser credential derived from a current expert-list entry and programme scope	May the adviser access a building-energy record or submit a funded-efficiency assessment?

The descriptions of the company register, BDEW services, dena expert list and ISO/IEC 42001 certification basis are supported by the respective authoritative sources [36–39]. None of those sources establishes a general W3C VC issuance service; credentialisation is the target pattern proposed here.

Each source domain retains its governance, liability, vocabulary, lifecycle and revocation rules. A target-domain verifier establishes trust by validating the required credentials, issuer authority, status and semantics for one purpose and time window. This approach preserves domain autonomy while allowing energy, manufacturing, mobility, construction and public-administration evidence to participate in a single authorisation decision.

Sector coupling provides the practical **transient cross-domain** rationale. A battery, heat pump, industrial plant or electric vehicle can participate simultaneously in energy, manufacturing, mobility and product-compliance data spaces. Manufacturing-X and DPP projects can reuse the legal-person root and add domain credentials for plant roles, asset provenance, conformity, market participation or regulatory access. AI agents then operate across these domains under a bounded mandate and locally enforced policy.

3 Comparative Results across China, Europe and the United States

The compared 2026 source families approach agentic trust from different organisational settings. They independently share several control patterns, although their terminology, scope and level of specification differ [1, 5–9]. The regional labels below identify each publication's organisational locus. They do not present any document as a national standard or government-endorsed regional architecture.

Source and locus	**Primary trust problem**	**Principal architecture patterns**	**Layer that remains open**
GSMA Greater China M-Trust	Autonomous agents interact across networks, organisations and trust domains	Multi-party evaluation, intent-aware authorisation, cross-domain trust propagation, heterogeneous trust roots and a collection–decision–enforcement loop	Credential profiles, legally grounded delegation and concrete PQC standards
United States B2C agentic-commerce initiatives: Google, Mastercard, Visa, Amazon and Apple	A merchant or payment provider must recognise an authorised shopping agent and verify what a natural person approved	Signed intent and cart mandates, identity–intent–action records, domain- and action-bound agent signatures, human checkout confirmation and platform-secured purchase actions	Common cross-platform credential semantics and status, legal effect of delegated intent, portability, revocation and consumer recourse
Anthropic in the United States	Enterprise agents can misuse legitimate permissions, poisoned context and connected tools	Cryptographically rooted identities, task-scoped permissions, memory protection, sandboxing and agent-oriented defensive operations	Portable legal-person identity, Business Wallets, W3C VC exchange and cross-company mandates
WE BUILD in the European Union	Agents require accountable authority for payments, contracts and critical-system actions	European Digital Identity, Business Wallets, verifiable credentials, cryptographically provable intent and mandates, mutual agent authentication and chains of proof	Detailed runtime trust evaluation, enforcement telemetry and PQC profiles
Manufacturing-X in Germany and Europe, energy data-X	Industrial data ecosystems require interoperable identity, trust, sovereignty and usage control across value chains	Federated Identity & Trust, data-space connectors, organisation and role credentials, policy negotiation and domain-specific rulebooks	Agent mandate semantics, common A2A profiles, evidence-graph profiles and cross-domain PQC migration
Spherity joint research	Material agent actions require both legal authority and current technical assurance	Legal-person identity (LPID), Power of Attorney (PoA) and AI Service Passport (AISP); Business and Agent Wallets; DIDs and VCs; verifier policy; enforcement gateways; receipts; PQC Corridors	Cross-jurisdiction conformance profiles and implementation evidence

The B2C initiatives make the human–agent branch visible in operational commerce. Google AP2 and Mastercard Verifiable Intent emphasise portable evidence of a consumer instruction; Visa emphasises agent recognition and domain- and operation-bound authorisation; Amazon and Apple rely more heavily on platform accounts, confirmation interfaces and controlled payment processing [28–33]. These mechanisms can support evidence of what a user approved. Their legal effect remains jurisdiction- and transaction-specific [35]. The B2B/B2G branch preserves the same identity–authority–action structure while replacing the sole natural-person root with a legal-person root, a representative or system delegation and organisational policy.

Anthropic's framework calls for cryptographically rooted identities, permissions scoped by task, protection of agent memory and an implementation workflow that covers identity, access scoping, sandboxing and input and output controls [5]. These principles align with M-Trust's identity-centric and intent-aware architecture. Anthropic focuses on enterprise deployment. Its published overview leaves legal-person identity, portable organisational credentials and cross-domain mandate semantics to an additional interoperability layer.

The WE BUILD non-paper supplies that layer. It treats verifiable credentials as the carrier for identities and signatures, requires cryptographically provable intent and mandates in agentic payments, and proposes mutual authentication of agents and the humans or organisations behind them [6]. Its use of the European Digital Identity Wallet (EUDIW) and Business Wallet frameworks gives the trust chain an authoritative institutional root. Carsten Stöcker of Spherity is one of the eight named co-authors.

Manufacturing-X supplies an industrial adoption environment rather than a competing identity standard. Its Identity & Trust framework supports secure, sovereign data exchange across sector data ecosystems, while Catena-X develops concrete DID, VC, participant-identity and credential profiles and is aligning them with the European Business Wallet [7–9]. The model therefore tests M-Trust's cross-domain premise in manufacturing supply chains, although present data-space exchanges are mainly connector- and service-oriented rather than optimised for autonomous A2A negotiation.

The document comparison yields six recurring controls: cryptographically verifiable identity; task-scoped authority; context-aware policy at action time; independently verifiable evidence and status; enforcement outside the agent's reasoning process; and reconstructable accountability evidence. M-Trust contributes the cross-domain evaluation loop. Anthropic contributes an enterprise security-control path. WE BUILD contributes institutionally anchored wallet and mandate semantics. Spherity's legal-person identity, Power of Attorney and AI Service Passport model composes these elements into an authority and assurance chain.

Figure 2 summarises this convergence at the level of B2B architecture patterns while retaining each source family's distinct scope and status.

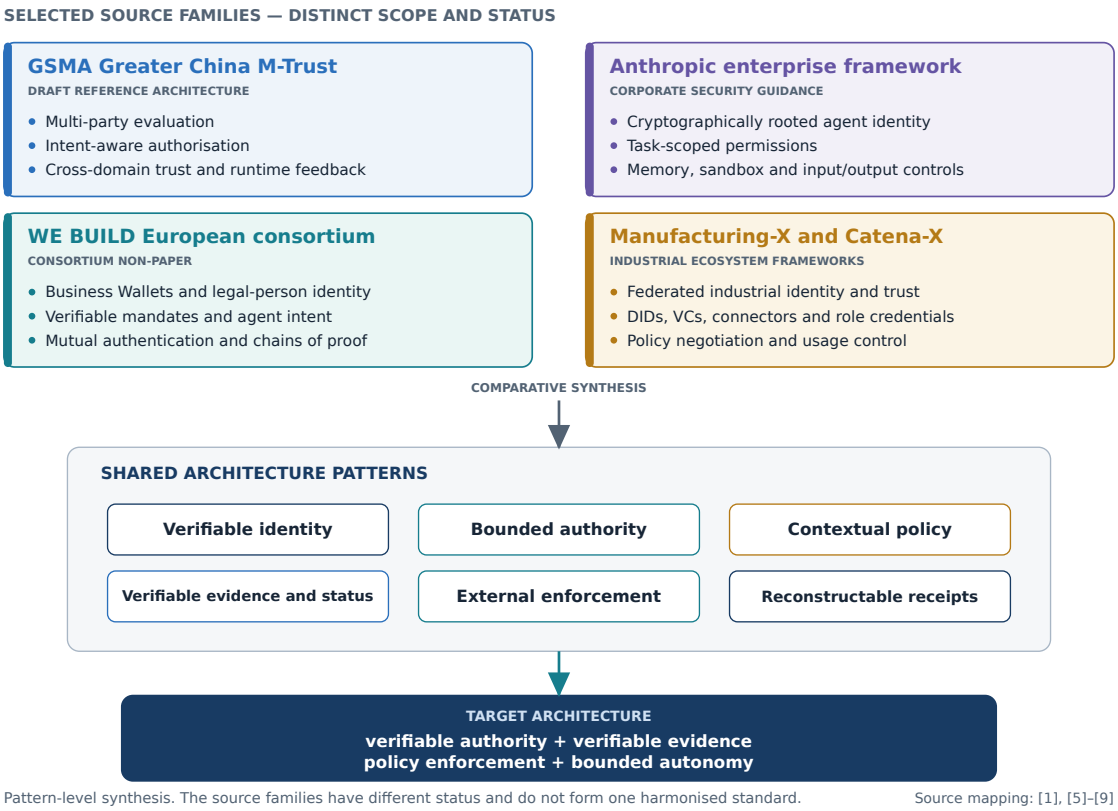

Figure 2: Converging agentic trust patterns for B2B use cases across selected 2026 sources. Source: Author's comparative synthesis based on [1, 5–9]. The organisational loci identify source settings and do not denote national standards.

4 Standards Requirements and Verifiable Trust Chains

The standards baseline determines whether M-Trust can carry authoritative evidence between trust domains. The draft explicitly names W3C Decentralized Identifiers, Post-Quantum Cryptography and Distributed Ledger Technology. It describes credential functions through generic terms. Its PQC treatment defines a direction rather than an implementable standards profile. At the evidence cut-off, the draft leaves the W3C Verifiable Credentials Data Model, enterprise wallet governance, legally grounded agent delegation and PQC algorithm profiles to subsequent work.

4.1 Documented Standards Coverage

The document supports the following classification:

- **W3C Decentralized Identifiers are explicit.** Section 5.1 describes the W3C DID standard as a core vehicle for distributed identity management [1, p. 19].
- **W3C Verifiable Credentials are absent as a named standard.** M-Trust uses generic terms such as "standardized trust credentials", "trust credentials", "identity credentials", "on-chain credentials" and "claims". It defines none of the W3C VC ecosystem roles or conformance requirements [1, pp. 9, 28].

- **Digital Wallets and Business Wallets are absent.** The paper describes credential management and exchange functions while leaving custody, presentation and organisational governance unallocated.
- **Agent Powers of Attorney are absent.** Intent-aware authorisation provides dynamic access logic while leaving legal-person delegation and representation semantics open.
- **PQC: explicit as an architectural direction.** The English and Chinese M-Trust drafts identify lattice-based cryptography, lightweight implementation and a hybrid architecture. Their reference lists cite NIST SP 800-207 for Zero Trust and contain no citation to NIST FIPS 203, FIPS 204 or FIPS 205. They contain no named Chinese PQC standard. The Chinese draft's reference to the SM series appears in the separate discussion of a software root of trust and describes classical national cryptography. The GSMA *Mobile AI* paper contains no PQC passage [1, 3].

Normative source — joint research: The target-architecture requirements in the following table are proposals of this article. The M-Trust coverage column reports the GSMA source position.

Component	M-Trust coverage	Requirement for the target architecture
W3C DID	Explicit	Identifier and key-control layer; method governance, resolver security, rotation, recovery and legal binding require profiles
W3C VC Data Model 2.0	Absent as a named standard; generic credential concepts appear	Carries portable identity, role, mandate, assurance and provenance evidence under governed profiles
Business Wallet	Absent; credential functions lack a holder-side governance component	Required organisational governance and automation layer for cross-company authority
Verifiable Presentation	Generic exchange and validation	Required for audience-bound, purpose-bound and privacy-preserving evidence exchange
Credential status and revocation	General revocation and updating	Required per credential, mandate, issuer, wallet and trust-anchor profile
Agent Power of Attorney	Intent and permission concepts	Required to bind a legal person to a specific agent, task and authority scope
PQC	Explicit at concept level; lattice basis, lightweight implementation and hybrid architecture named; no NIST or Chinese PQC standard cited	Algorithms, parameter sets, protocols, proof formats, lifecycle and migration governance require profiles
DLT/blockchain	Explicit and optional	Appropriate for selected shared-registry and integrity functions under explicit governance

4.2 Standards Comparison: M-Trust and Manufacturing-X

M-Trust and Manufacturing-X operate at different architectural levels. M-Trust is a draft conceptual trust architecture for agentic telecom networks. Manufacturing-X is a policy and implementation umbrella for federated industrial data ecosystems, with specifications distributed

across initiatives such as Catena-X. The comparison therefore concerns functional coverage rather than equivalence between two standards [1, 7–9].

Standards dimension	**M-Trust**	**Manufacturing-X and Catena-X implementation path**	**Convergent requirement**
Primary scope	Agentic 6G/network trust	Sovereign industrial data exchange across value chains	Cross-organisational, machine-actionable trust
Identity model	User, device and agent identities; W3C DIDs named	Participant identity, Business Partner Number, DIDs and federated identity services; Catena-X CX-0049 and CX-0149 provide DID-document and data-space identification profiles	Bind technical actors to accountable legal persons
Credential model	Generic trust and identity credentials	W3C VC-oriented company, membership and framework credentials; Catena-X CX-0050 profiles the Framework Agreement Credential	Govern issuer, holder, verifier, schema, status and holder-binding semantics
Trust domains	Multi-party and cross-domain trust with heterogeneous roots	Domain rulebooks and federated anchors across Manufacturing-X ecosystems	Preserve issuer provenance and verifier sovereignty across domains
Authorisation	Intent-aware, dynamic permissions and task context	Contract, policy and usage-control negotiation at connector and service layers	Express task, purpose, audience, scope, time and delegation depth
Exchange protocols	Future 6G authentication, A2A security and MCP Security concepts	Dataspace Protocol, Decentralized Claims Protocol, connector APIs and credential exchange	Bind message, contract and data transfer to the same authority context
Wallet role	Unspecified	Emerging alignment with European Business Wallets and credential services	Business Wallet as legal-person and delegation control plane
Trust registry/VDR	DLT optional for identity, permissions, reputation and audit	Federated identity services, trust anchors, credential status and ecosystem registries	Resolve authoritative issuers, semantics and status across domains
Agentic assurance	Behaviour, reputation, intent and runtime feedback	Trusted AI work is emerging through Catena-X governance and AI Service KIT planning	Add AI Service Passports, runtime attestation and evidence graphs
PQC	Directional: lattice-based, lightweight and hybrid	No common Manufacturing-X PQC profile identified in the reviewed framework	Define cross-ecosystem cryptographic profiles and migration corridors

This mapping suggests a division of labour. M-Trust contributes the continuous agentic evaluation loop and telecom trust anchors. Manufacturing-X contributes industrial governance, data-sovereignty and connector ecosystems. W3C VCs, DIDs, Business Wallets, DCP/DSP profiles and governed VDRs can form their shared interoperability layer.

4.3 Functions in a Verifiable Trust Chain

Each standard and component performs a distinct function. A DID identifies a subject and enables a controller to prove control through the associated DID document. A VC carries issuer-authored claims about one or more subjects. A Verifiable Presentation can convey selected evidence and, under an appropriate securing mechanism, bind the presentation to a holder, verifier or challenge. A Business Wallet governs keys, credentials, presentations and delegated actions for a legal person. Trust registries identify accepted issuers and credential profiles. A policy engine decides whether the combined evidence authorises the requested action. A signed receipt records the decision and execution state.

Together, these functions establish the following verifiable chain:

authoritative source → legal person → authorised representative → agent or workload → task mandate → verifier policy → controlled action → signed receipt

This chain supplies legal and evidentiary semantics to M-Trust. The later sections specify its organisational, semantic, cryptographic and registry components.

5 Organisational Identity and Verifiable Authority

This section defines the identity and delegation components that connect an agent to an accountable organisation and a bounded mandate.

5.1 Business Wallets as the Organisational Control Plane

A Business Wallet represents a legal person and governs the credentials, keys, representatives, system identities, branches, approval rules and liability boundaries associated with that organisation. Its function therefore extends beyond credential storage.

Normative source — joint research: The following Business Wallet requirements extend M-Trust with Spherity's organisational identity and authority model. They are proposals of this article.

An enterprise-grade Business Wallet SHOULD provide:

- register-backed legal-person identity and identifier matching;
- roles for directors, representatives, employees, services, devices and agents;
- issuance, receipt, storage and presentation of credentials;
- electronic signatures, seals, timestamps and secure document exchange;
- powers of attorney with scope, purpose, audience, value, geography and time constraints;
- multi-user approval, separation of duties and threshold controls;
- credential, mandate, wallet-unit, issuer and trust-anchor status verification;
- protected key custody, recovery, rotation and incident response;
- policy application programming interfaces (APIs) for machine-to-machine and agent-to-agent transactions;
- signed action receipts and lifecycle audit evidence; and

- cryptographic agility across classical, hybrid and post-quantum profiles.

The European Business Wallet proposal provides a prospective legal and cross-border framework for several of these functions. The Commission proposal includes wallet-owner identification data, automated wallet interaction, auditable and revocable role mappings, electronic seals and support for trusted cross-border business processes [7]. The Council's June 2026 position covers digital identity, trusted documents, electronic signatures, timestamps and seals, delegation of legal authority and secure business-to-business and business-to-government communication [40]. It also states that the proposed authorisation system leaves powers of attorney and legal mandates under national and Union law unaffected. The legislative procedure remained open at that stage. The European Economic and Social Committee calls for harmonised identity matching, harmonised mandate structures, a clear distinction between personal and business wallets and international operability [41].

5.1.1 Stable Business Wallet and Bounded Agent Wallet

The target architecture separates two wallet scopes.

Normative source — joint research: The wallet separation below is proposed by this article.

Wallet scope	Primary function	Typical contents	Lifecycle
Business Wallet	Represents the legal person and governs delegation	Legal-person identity, company roles, licences, qualified certificates, trust relationships, delegation policy	Stable, governed and recoverable
Agent Wallet or wallet unit	Represents a specific agent, service or workload	Agent identifier, workload attestation, short-lived PoA, task capability, AI Service Passport reference	Short-lived, scoped, rotatable and revocable

The Business Wallet remains the root of organisational authority. The Agent Wallet receives the minimum authority required for a defined task. Runtime systems enforce the mandate at every material action.

This pattern supports regional legal variation. Europe can anchor identity in electronic identification, authentication and trust services (eIDAS), national registers and the emerging European Business Wallet framework. North American and Asia-Pacific deployments can use local authoritative registers, regulated issuers, Legal Entity Identifier (LEI) or verifiable LEI (vLEI) evidence and sector trust frameworks. W3C data models and open exchange protocols can provide an interoperability layer across these legal roots.

5.2 Decentralized Identifiers and Identifier Governance

M-Trust describes DIDs as a route beyond dependence on traditional central certificate authorities [1, p. 19]. W3C DID Core defines a broader technology-neutral bridge: DID methods can use centralised, federated or decentralised identifier systems [10]. This distinction matters because global agentic networks will combine Web PKI, telecom identities, government registers, permissioned infrastructures and distributed registries.

A DID document can expose verification methods, authentication relationships, capability-invocation keys and service endpoints. Its primary proof concerns control of the identifier. Legal existence, representation rights, licences and authority arise through credentials from recognised issuers.

Normative source — joint research: The following DID profile is developed in this article. M-Trust supplies the DID direction and cross-domain trust objective [1].

The following normative profile follows:

1. A verifier **MUST** distinguish DID control, credential subject identity and legal-person authority.

2. A verifier **MUST** resolve the DID under the rules and current state of the selected DID method.

3. A trust framework **MUST** define accepted DID methods, resolver requirements, update semantics, monitoring and recovery.

4. A wallet **MUST** support key rotation, compromise response and cryptographic migration.

5. An agent DID **MUST** bind to an authoritative organisational credential and a current delegation before a material cross-company action.

6. A DID document **MUST NOT** carry confidential business data, personal data, agent memory or transaction content.

7. A privacy profile **SHOULD** use pairwise or purpose-specific identifiers where persistent correlation creates risk.

The choice of DID method is therefore a governance decision. A ledger-based method adds shared consensus and operational dependencies. A web-based method inherits Domain Name System (DNS) and Web Public Key Infrastructure (Web PKI) dependencies. An ephemeral method favours local control and limited lifecycle functions. Risk, jurisdiction, privacy, recovery and evidence horizon determine the profile.

5.3 Verifiable Credentials as Semantic Trust Objects

The W3C VC Data Model 2.0 defines issuer, holder, verifier and verifiable data registry roles [11]. It gives M-Trust a standard way to express tamper-evident, machine-readable credentials and presentations. Cryptographic verification can establish authorship and integrity. The VC Data Model does not provide a complete authorisation framework; verifier policy must establish whether the issuer, claims and evidence are fit for a specific purpose [11].

This is essential for M-Trust. Cross-domain trust propagation becomes an exchange of evidence into a new decision domain. The receiving verifier evaluates the evidence under its own trust framework, risk appetite and policy.

5.3.1 Credential Profiles for Trusted Agentic AI

Normative source — joint research: The following credential profiles are proposed by this article as the semantic layer that M-Trust leaves open.

Credential profile	Authoritative issuer	Subject or holder	Decision purpose
Legal Person Identity Credential	Business register, qualified trust service or recognised identity authority	Organisation and Business Wallet	Proves the accountable organisation
Representation or Company Role Credential	Organisation or authoritative register	Director, employee, service operator or wallet unit	Proves capacity to act for the organisation

Credential profile	Authoritative issuer	Subject or holder	Decision purpose
Market Role or Licence Credential	Regulator, scheme authority or governed sector issuer	Organisation, branch or service	Proves regulated or ecosystem capacity
Agent Identity Credential	Organisation, wallet provider or workload identity authority	Agent instance or agent service	Binds the technical actor to its operator
Agent Power of Attorney Credential	Authorised organisational wallet	Agent Wallet or workload	Defines task, action, audience, purpose, value and time scope
AI Service Passport (AISP)	Provider, operator, assessor and monitoring services under governed profiles	AI service or deployed agent	Exposes provenance, version, intended use, risk, testing, evaluation, verification and validation (TEVV) evidence and current status
Runtime Attestation	Trusted execution, platform or attestation service	Running workload and environment	Proves current technical state
Action Receipt	Policy and enforcement gateway, optionally counter-signed by participants	Transaction or action	Binds authority evidence, policy decision, execution and time

Each profile requires semantic governance. Schema identifiers, controlled vocabularies, issuer accreditation, status mechanisms, proof suites, holder binding, privacy rules and conformance tests determine interoperability.

5.3.2 Credential Verification, Validation and Authorisation

Credential verification and business validation form separate steps:

- **Verification:** proof integrity, issuer identifier, validity period, status, schema conformance, presentation challenge and holder binding where the profile requires it.
- **Trust resolution:** issuer authority, trust list, accreditation, jurisdiction and accepted credential profile.
- **Validation:** fitness of the claims for the requested action under local business rules.
- **Authorisation:** effective mandate, context, risk, runtime state and policy decision.
- **Evidence retention:** signed receipt with the status and policy snapshot used at decision time.

Normative source — joint research: The following rule combines the W3C VC trust model with Spherity's verifier-sovereign authorisation model.

M-Trust's dynamic trust chains SHOULD transfer evidence and provenance, with authority remaining under the receiving verifier's policy. A verifier **MUST NOT** convert reputation, endorsement or an upstream trust score into transitive legal authorisation. This rule applies the VC Data Model's separation between cryptographic verification and the accompanying authorisation framework [11].

5.4 Verifiable Authorisation Chains and Assurance Graphs

The minimum authorisation path for a material agent action is:

authoritative source → Business Wallet/legal person → representative or role → Agent Wallet/workload → task mandate → local policy decision → executed action → signed receipt

The path establishes provenance of authority. Its effective scope narrows at each delegation step. Let $A(C_i)$ denote the set of actions permitted by credential C_i, $A_P(X_t)$ the actions allowed by local policy in context X_t, and $A_R(t)$ the actions allowed by the current runtime state. Effective authority at time t is:

$$EA_t = \left(\bigcap_{i=1}^{n} A(C_i) \right) \cap A_P(X_t) \cap A_R(t)$$

Let V_t represent the conjunction of the required identity, issuer, status, assurance and cryptographic predicates. The enforcement point permits action a only when:

$$Permit(a, t) = 1 \quad \Leftrightarrow \quad a \in EA_t \wedge V_t = 1$$

This model prevents a broad permission at one layer from expanding a narrower delegation at another layer.

5.4.1 Linear Authority Chains and Contextual Assurance Graphs

The word "chain" describes delegation lineage. The complete assurance structure is a graph because evidence comes from multiple independent sources:

- a register or trust service establishes legal-person identity;
- the organisation establishes representation and agent delegation;
- a sector authority establishes a licence or market role;
- an assessor establishes TEVV or conformity evidence;
- a platform attester establishes runtime state;
- a data-space governance authority establishes participation and policy context;
- product and data issuers establish provenance and status; and
- the local verifier establishes the final authorisation decision.

M-Trust's Multi-party trust dimension provides the natural home for this assurance graph. W3C VCs provide portable evidence objects. Spherity's Verifiable Linked Knowledge Graph supplies provenance, validation, freshness, uncertainty and traversable evidence paths [34]. The Business Wallet supplies accountable authority [25].

5.4.2 Normative Authorisation Requirements

Normative source — joint research: The following requirements are the authorisation-chain profile proposed by this article. M-Trust contributes continuous evaluation, intent awareness and cross-domain propagation [1].

For every material cross-company action, the verifier:

1. **MUST** authenticate the agent or workload and bind it to the presented holder proof.

2. **MUST** identify the accountable legal person through an accepted authoritative credential.

3. **MUST** verify every delegation step and the delegator's authority to delegate.

4. **MUST** compute the intersection of scopes across the chain.

5. **MUST** evaluate audience, purpose, action, resource, time, value, geography and separation-of-duties constraints.

6. **MUST** resolve credential, mandate, wallet, issuer and trust-anchor status at decision time.

7. **MUST** apply local policy as the final authority predicate.

8. **MUST** bind the decision to a transaction identifier, nonce and replay boundary.

9. **MUST** create a signed receipt that links evidence, policy version, status snapshot, execution result and timestamp.

10. **MUST NOT** treat an identity proof, DID, trust score, model score or transport session as sufficient legal authority by itself.

6 Cryptographic and Registry Infrastructure

This section specifies cryptographic migration and shared-registry functions for long-lived, cross-domain evidence.

6.1 Post-Quantum Cryptography and Cryptographic Agility

GSMA source position: M-Trust places PQC in the trust foundation. The English draft states that quantum computing threatens traditional cryptography through Shor's and Grover's algorithms. It identifies long-term confidentiality as a design concern, describes lattice-based cryptography as a promising family, calls for lightweight implementations on constrained agents and anticipates a hybrid architecture in which symmetric, classical public-key and post-quantum cryptography coexist. It expects PQC to support encryption, signatures and agent identity authentication [1, pp. 27–28, 30]. The Chinese draft contains the same substantive claims [1].

The source wording requires one technical qualification. M-Trust uses "asymmetric encryption" as an umbrella for encrypted communication, digital signatures and identity authentication. Standards profiles need separate primitive classes. A key-encapsulation mechanism establishes a shared secret, while a digital-signature algorithm authenticates signed data and its signatory. The NIST standards make this separation explicit [13–15].

6.1.1 PQC Claims and Standards Attribution

The GSMA source corpus establishes a PQC direction without selecting a standard. This distinction matters because an algorithm family provides neither wire-level interoperability nor

lifecycle governance.

Source	PQC treatment	Named PQC algorithm or standard	Finding
M-Trust English and Chinese drafts	Explicit; quantum threat, lattice-based cryptography, lightweight implementation and hybrid architecture	None	Architectural direction without algorithm identifiers, parameter sets, protocol profiles, migration dates or conformance references
M-Trust reference list	NIST appears through SP 800-207 for Zero Trust	No FIPS 203, FIPS 204 or FIPS 205 citation	The NIST Zero Trust citation provides no PQC profile
M-Trust Chinese and English classical-cryptography passage	The SM series and lightweight elliptic-curve cryptography support a software-based root of trust	No named Chinese PQC algorithm or PQC standard	The SM-series reference concerns established classical Chinese national cryptography and is separate from the PQC passage
GSMA Mobile AI paper	No PQC treatment	None	The document offers no additional PQC standard selection [3]
This article's implementation profile	Separates key establishment from signature functions and adds lifecycle migration	NIST FIPS 203, FIPS 204 and FIPS 205 [13–15]	Joint-research proposal for interoperable PQC Corridors

The result is unambiguous for the reviewed editions: the M-Trust reference list cites neither the NIST PQC standards nor a named Chinese PQC standard. It cites NIST SP 800-207 for the Zero Trust baseline. The text uses PQC as a technology class and lattice-based cryptography as an algorithm family. The NIST FIPS mapping below originates in this article and the cited Spherity research [13–15, 42, 43].

Normative source — joint research: The following functional allocation turns the GSMA direction into a standards-based implementation profile.

1. **ML-KEM (FIPS 203)** supports quantum-resistant shared-secret establishment [13].

2. **ML-DSA (FIPS 204)** supports quantum-resistant digital signatures based on module lattices [14].

3. **SLH-DSA (FIPS 205)** provides a stateless hash-based signature alternative [15].

Algorithm selection remains a profile decision based on security level, ciphertext and signature size, latency, constrained-device capability, hardware security module support, implementation maturity and evidence horizon. Selection at the primitive layer remains insufficient for interoperability: protocols, proof formats, key representations, algorithm negotiation, downgrade controls and conformance tests also require profiles.

6.1.2 Regional PQC Standardisation and Research Trajectories

The regional comparison requires separation between published standards, migration policy and research direction. "Western PQC" has no single profile: the United States supplies the dominant final algorithm standards, the European Union coordinates a common transition, Germany applies additional national assurance and implementation criteria, and China is conducting an indigenous commercial-cryptography selection programme [13–19, 44].

Region	Published preference or programme at the evidence cut-off	Architectural consequence
United States	NIST has standardised ML-KEM, ML-DSA and SLH-DSA in FIPS 203–205. Federal migration guidance and the National Security Agency's CNSA 2.0 profile give lattice-based algorithms the principal operational role, with stateful hash-based signatures retained for selected use cases.	Global products increasingly inherit NIST identifiers, parameter sets and test vectors; regulated profiles still require protocol integration, validated implementations and downgrade controls.
European Union	The coordinated PQC roadmap calls for Member States to begin transition by the end of 2026 and prioritises high-risk systems and critical infrastructure for earlier migration. It draws heavily on internationally standardised algorithms while preserving European risk governance, implementation assurance and crypto-agility.	Cross-border identity, wallet, data-space and trust-service profiles need a coordinated transition schedule and long-term validation rules rather than isolated algorithm upgrades.
Germany	BSI technical guidance now includes the final NIST PQC standards and evaluates deployment through national technical guidelines, implementation-security work and migration research. Germany's emphasis is diversity, high-assurance implementation and critical-infrastructure readiness rather than a separate incompatible wire standard.	Business Wallets, qualified trust services, industrial controllers and VDR nodes require BSI-aligned implementation profiles, hardware support and conformance evidence.
China	The Institute of Commercial Cryptography Standards launched its Next-Generation Commercial Cryptographic Algorithms programme in 2025 and continued algorithm and implementation-evaluation work through 2026. M-Trust characterises lattice-based cryptography as promising and separately cites the classical SM suite for software roots of trust. The reviewed M-Trust editions name no selected Chinese PQC algorithm or national PQC standard.	Interoperability planning should anticipate an indigenous Chinese profile while avoiding premature assumptions about its algorithm identifiers or final status. Cross-border corridors need explicit negotiation and equivalence testing.

The strategic distinction is timing and governance. The United States has final algorithm standards; Europe and Germany are operationalising migration through coordinated and national assurance frameworks; China is developing a domestic next-generation selection and standardisation route. M-Trust's generic hybrid architecture can accommodate these trajectories, yet it supplies no algorithm-agility or cross-profile negotiation mechanism. That missing layer is material for multi-polar agentic networks.

6.1.3 Public Key Trust Fabric

PQC readiness covers every component that relies on public-key authenticity, confidentiality or long-term evidence:

Normative source — joint research: The following end-to-end migration scope extends M-Trust's trust-foundation concept.

Trust-fabric component	**PQC requirement**
Issuers and trust services	Hybrid/PQC signing profiles, protected keys, rollover and reissuance
Business and Agent Wallets	Crypto-agile key stores, holder binding, recovery and device/hardware security module (HSM) support
Verifiers and policy engines	Algorithm negotiation, allow-lists, downgrade resistance and policy versioning
DIDs and DID methods	PQC verification methods, update rules, historical resolution and key transition
VCs and Verifiable Presentations	PQC-capable proof formats, selective disclosure strategy and interoperability tests
Trust and status lists	Resilient signatures, availability, emergency rollover and offline validation
Verifiable data registries and DLT	PQC-protected governance keys, state transitions, consensus identities and checkpoints
Secure exchange	Hybrid/PQC key establishment and authenticated channels
AI Service Passports (AISPs), Digital Product Passports (DPPs) and evidence graphs	Signature renewal, timestamps, provenance continuity and archival validation
Action receipts	Long-term validation data, renewal and evidentiary preservation

6.1.4 PQC Corridors

Spherity Research proposes a **PQC Corridor** as a bounded migration domain [42,43]. A corridor fixes the use case, actors, legal assurance, evidence horizon, trust boundary, cryptographic profile and operational controls. This approach addresses the network effect: an issuer, wallet or verifier gains operational PQC value when counterparties recognise compatible profiles and lifecycle rules.

Normative source — joint research: The following corridor requirements originate in Spherity research and extend the GSMA paper's generic hybrid-PQC direction.

Regulators and operators SHOULD establish corridors for high-value flows such as cross-border organisational onboarding, data-space access, Digital Product Passport signing, AI-agent delegation and regulated agent-to-agent transactions. Each corridor SHOULD produce:

- a Cryptographic Bill of Materials and Trust Anchor Bill of Materials;
- accepted classical, hybrid and PQC profiles;
- issuer, wallet, verifier, registry and status interoperability tests;
- downgrade, deprecation and emergency-rollover rules;
- reissuance and long-term validation procedures;

- telemetry, incident and degraded-mode runbooks; and
- procurement clauses and auditable evidence.

6.2 Distributed Ledger Technology and Bounded Registries

M-Trust presents DLT as one implementation option for distributed trust evaluation, evidence integrity, credential management, reputation, permissions and audit. It expressly classifies blockchain as one possible implementation rather than a mandatory M-Trust component [1, p. 28]. This design freedom aligns with the W3C VC model, which recognises trusted databases, government identity databases, decentralised databases and distributed ledgers as possible verifiable data registries [11].

The architectural rule is straightforward:

Consensus establishes agreement on ledger state. Issuer governance establishes claim authority. Verifier policy establishes fitness for purpose.

6.2.1 Suitable Shared Registry Functions

Normative source — joint research: M-Trust presents DLT as optional [1]. The following list defines the bounded functions proposed by this article.

DLT MAY support:

- DID operations for selected DID methods;
- issuer and trust-anchor metadata;
- credential status or revocation commitments;
- schema and policy-version integrity anchors;
- timestamps and ordering of critical state changes;
- hashes of action receipts and audit batches;
- consortium governance decisions; and
- multi-party settlement or spectrum-allocation state where shared execution is required.

6.2.2 Off Ledger Evidence

Normative source — joint research: The following data-minimisation and governance requirements are proposals of this article.

The architecture **MUST NOT** place full VCs, personal data, confidential business attributes, prompts, agent memory, raw behavioural telemetry, proprietary policy or detailed Digital Product Passport or Digital Battery Passport records on a broadly replicated ledger. Wallets, data-space connectors and governed repositories retain these data. The ledger stores the minimum shared state or cryptographic commitment required for verification.

Every DLT profile **MUST** define governance membership, validator authority, jurisdiction, privacy, data retention, fork handling, upgrade rules, availability, finality, key compromise, PQC migration and exit procedures. A conventional registry remains suitable where one accountable authority already owns the source of truth.

6.2.3 EBSI, Qualified Electronic Ledgers and Cross-Domain VDRs

EBSI provides a European shared trust-layer pattern for verifiable credentials and registry-backed verification [20]. Its relevance to M-Trust lies in replicated, governed state. An organisation that operates a conformant node can validate ledger state from its local copy and avoid a live query to one commercial registry intermediary for every decision. This improves availability, auditability and evidentiary continuity across trust domains.

The legal term **qualified electronic ledger** requires precision. Regulation (EU) 2024/1183 defines an electronic ledger and gives data recorded in a qualified electronic ledger a presumption concerning unique and accurate chronological ordering and integrity. Qualification depends on provision by one or more qualified trust service providers and compliance with Article 45l [21]. EBSI is therefore a potential infrastructure substrate for a qualified-ledger service; EBSI participation alone establishes neither qualified status nor the legal presumption. A concrete service requires the applicable provider qualification, supervision, technical profile and conformity evidence.

Normative source — joint research: The following EBSI/VDR profile extends M-Trust and Zero Trust with a bounded European verification backbone.

A cross-domain VDR SHOULD make the following minimal states independently resolvable:

- DID documents or method-specific identifier state;
- trust-registry entries linking trust domains, accredited issuers and accepted credential profiles;
- privacy-preserving credential-status or revocation entries containing no personal data for regulated-industry use;
- semantic schema, controlled-vocabulary and policy-version identifiers or integrity commitments; and
- time-ordering and integrity commitments for material governance and receipt events.

Local replication reduces online dependency on an external lookup service. It preserves the need to evaluate issuer authority, governance membership, node software, consensus finality, legal qualification and the authenticity of off-ledger evidence. This distinction is central to both Zero Trust and M-Trust: the VDR supplies current, tamper-evident verification inputs, while the local policy decision point determines whether those inputs authorise the requested action.

Cross-domain trust registries are especially important for transient trust. They allow an energy verifier, for example, to discover an accepted company-register issuer, a BDEW market-role issuer and a TIC certification issuer without collapsing those governance regimes into a single root. The VDR records the relationship and status needed for verification; each source domain retains authority over its own claims.

7 Applications in Data Spaces, Digital Product Passports and Trusted AI

7.1 Data-Space Identity and Protocol Layers

The Data Spaces Support Centre Blueprint recommends W3C VC 2.0 and DIDs for identity and attestation management. It identifies OpenID for Verifiable Credentials and the Eclipse

Dataspace Decentralized Claims Protocol as complementary credential-exchange approaches [4]. The Blueprint also separates identity and attestation management, trust frameworks and access-policy enforcement. This separation creates a direct bridge from M-Trust's cross-domain trust objective to European data-space architecture.

For restricted data, authentication answers **which organisation is requesting access**. A governed role credential answers **in which capacity it acts**. A Power of Attorney answers **which agent may act for that organisation**. A policy engine answers **which operation the evidence permits for this asset and purpose**.

Protocol choice follows function:

Protocol family	**Primary function**	**M-Trust relevance**
DIDComm Messaging	Authenticated, encrypted, transport-independent messages whose relationships can be rooted in DIDs [45]	Suitable for wallet and agent exchanges that require pairwise secure messaging and correlation control
Dataspace Protocol (DSP)	Interoperable catalog, contract-negotiation and transfer processes under usage control [46]	Carries data-space business negotiation and transfer context across organisations
Decentralized Claims Protocol (DCP)	Presentation and exchange of identity and entitlement claims for data-space participants [47]	Supplies credential evidence to connector-side identity and policy decisions
OpenID4VCI/OpenID4VP	Credential issuance and presentation using widely deployed web patterns	Connects Business Wallets, issuers and verifiers to credential lifecycle operations
A2A and MCP security profiles	Agent collaboration and tool or resource invocation	Bind delegated authority, task context and evidence to agent and tool actions

DSP and DCP are optimised for governed machine-to-machine data-space interaction, while DIDComm provides a general secure message layer and emerging A2A protocols address agent coordination. Their convergence point is the authorisation context: identity, mandate, policy, contract and data-transfer state should refer to the same transaction and evidence set.

The energy data-X reference case demonstrates the core pattern for federated energy-data exchange: authentic sources and governed issuers supply company and market-role evidence; Business Wallet concepts manage the evidence; and a connector-side policy engine produces an access decision [12]. The reported implementation covers organisational and market-role evidence in an energy-data use case. Its transfer to differentiated access for DPPs and Digital Battery Passports (DBPs) is an architectural application rather than a completed implementation claim [12].

Trusted AI adds a further evidence layer. An agent's mandate establishes authority. Its AISP, TEVV results, runtime attestation and evidence graph establish whether the relevant service version, operational state and supporting data satisfy the relying party's assurance policy.

7.2 Manufacturing-X and the Catena-X Trusted AI Path

Manufacturing-X applies federated identity and trust to industrial supply chains. Catena-X offers the most mature sector implementation environment: its identity work uses DIDs and W3C VCs, evolves from a single anchor towards federated trust and aligns its credential framework with

the European Business Wallet [8, 9]. Its connectivity and credential services bind participant identity to connector-mediated catalog, contract and data-transfer processes.

The Catena-X Trusted AI Expert Group was established in 2026 under its Architecture Management Committee. Its published objectives include the registration, discovery, contracting and governance of cross-company AI services, with planned policy, semantic and conformity elements in an AI Service KIT and support for distributed and multi-agent scenarios [9]. At the evidence cut-off, these are programme objectives and planned specifications rather than completed standards or deployment results.

Catena-X therefore provides a useful real-world test environment for an agentic trust model. DSP, DCP and connector APIs originate in M2M data sharing and are less specialised for high-frequency A2A collaboration than protocols designed expressly for agents. When AI agents act through these interfaces, however, the ecosystem can test whether legal-person credentials, membership and role evidence, policy negotiation, usage control and receipts remain coherent across autonomous transactions. The AI Service Passport and evidence graph proposed here extend that established trust model with current AI-system assurance.

7.3 IPCEI-AI Outlook: From Control Plane to Evidence Data Plane

The emerging Important Project of Common European Interest for Artificial Intelligence (IPCEI-AI) is intended to mobilise cross-border industrial collaboration around next-generation AI, compute, data and sector adoption. Under the EU IPCEI framework, participating projects must form an integrated multi-state undertaking, address substantial technological or market risk and create positive spillovers beyond the funded participants [48, 49]. At the evidence cut-off, IPCEI-AI remains in preparation and candidate-project formation; it is therefore an opportunity for architectural alignment rather than evidence of an implemented common infrastructure.

Its ecosystem power lies in linking research organisations, cloud and compute providers, data-space operators, model and agent developers, industrial adopters, conformity-assessment bodies and public authorities. A common M-Trust-like control plane can provide organisational identity, agent delegation, trust-domain discovery, policy enforcement and action receipts across that value chain. The next phase can extend the same verifiability into the data plane through a **verifiable evidence graph**.

A verifiable evidence graph is a machine-resolvable graph whose nodes represent evidence objects—data provenance, model and agent versions, licences, DPP claims, test results, ISO/IEC 42001 certificates, runtime attestations, incidents, policy decisions and action receipts—and whose edges carry signed provenance and dependency relations [34]. It matters because a final AI output rarely proves its own lineage, legal basis, safety state or fitness for purpose. The graph lets a verifier traverse from an action to the accountable organisation, mandate, software version, source evidence, assessment and status used at decision time.

Normative source — joint research: The following IPCEI-AI direction is a proposal of this article.

IPCEI-AI candidate projects SHOULD adopt interoperable Business Wallet, VC, DID, VDR, policy and evidence-graph profiles as shared infrastructure. Early conformance work can convert European legal and industrial assets into a deployment advantage; delayed alignment would force projects to reconcile incompatible identity, authority and evidence models after deployment. European implementation therefore requires speed, reusable reference components and cross-project testing alongside careful legal and security assurance.

8 Alignment with the Spherity Trusted AI Architecture

Spherity separates three cumulative functions [26, 27]:

Layer	Core question	Principal mechanisms
AI Governance	Is the deployment legitimate and accountable?	Law, governance, risk classification, oversight, audit and remedies
Trustworthy AI	Does the system exhibit the required qualities?	Safety, robustness, security, fairness, transparency, privacy, TEVV and lifecycle monitoring
Trusted AI	Can the actor, authority, evidence, status and action be verified now?	Legal-person identity, agent identity, PoA, AISP, provenance, status, policy, attestation and receipts

The Spherity *Trusted Industrial AI* deck expresses the operational chain as **LPID** $\rightarrow$ **PoA** $\rightarrow$ **AISP** [27]:

1. **LPID:** register-backed legal-person identity held in the Business Wallet. A European profile can combine the EU Company Certificate (EUCC, used here as shorthand), the European Unique Identifier (EUID) used in the Business Registers Interconnection System and European Business Wallet owner identification data (EBW-OID, used here as shorthand). A verifiable Legal Entity Identifier (vLEI) credential supplies a complementary, globally usable organisational and role identity for cross-jurisdiction transactions [7, 50, 51];

2. **PoA:** revocable delegation from the company to a specific agent with defined scope and validity; and

3. **AISP:** the AI Service Passport containing operator, provenance, model or agent version, intended use, TEVV, runtime controls and status.

The identifiers serve complementary trust functions. EUCC, EUID and EBW-OID connect the wallet to European company-law and register infrastructure. LEI and vLEI connect the same organisation and its authorised roles to a globally recognised entity identifier. Identifier matching across these sources reduces duplicate organisational identities and lets a verifier select the legal and assurance root appropriate to its jurisdiction and transaction.

This chain extends Zero Trust from technical identity and resource access to legal authority and current assurance. The mapping to M-Trust is an analytical result of this article:

M-Trust pillar	Spherity-aligned implementation
Multi-party trust	Authoritative issuers, trust lists, verifiable data registries, evidence graphs, TEVV providers and local risk inputs
Agentic Intent-aware trust	PoA scope, task purpose, AISP, runtime context, policy engine and bounded action capability
Cross-domain trust	Business and Agent Wallets, W3C VCs/VPs, DIDs, status, trust frameworks, data-space connectors and action receipts

8.1 Work with Bundesanzeiger Verlag and the German Company Register

The joint research programme described by the author includes ongoing work with Bundesanzeiger Verlag around German company-register identity. The German Company Register is the central

platform for company data and provides an authoritative source from which an accepted issuer can derive legal-person and representation evidence [38]. This programme statement describes current collaboration direction rather than a public production-deployment claim.

The target integration uses register evidence at issuance and status events, then allows the Business Wallet to present a purpose-bound LPID credential to data spaces, public authorities, supply-chain partners and AI-agent verifiers. The register remains the authoritative source. The credential provides portability and automation; the verifier resolves current issuer and credential status and applies its own policy. This separation avoids turning a wallet self-assertion into company identity and supports a common legal-person root across Energy Data-X, Manufacturing-X, DPP and Trusted-AI domains.

8.2 Alignment with GLEIF for Cross-Jurisdiction vLEI Trust

The programme also pursues ongoing alignment with the Global Legal Entity Identifier Foundation (GLEIF) on vLEI credentials for cross-jurisdiction company and role identity. This is an author-reported collaboration direction. GLEIF's 2026 working paper on agentic AI in payments analyses computational verification of the organisation behind an agent and its delegated authority; the paper expressly presents research rather than an official institutional position [50]. GLEIF's European Business Wallet analysis argues that vLEI can connect company and representative identity across European and global transactions [51].

Within the proposed architecture, a vLEI credential complements European register credentials rather than replacing them. The Business Wallet can present EUCC/EUID/EBW-OID evidence for European legal effects and vLEI evidence for global entity matching and authorised organisational roles. A cross-jurisdiction verifier evaluates both under its local trust framework, creating an interoperable chain from authoritative company identity through a role and PoA to the agent action.

The joint research programme described here pursues this target architecture with partners in Europe, North America and Asia-Pacific and through W3C-related standardisation work. This statement records the author's programme position. Independent validation requires published profiles, reference implementations, conformance results and cross-jurisdiction pilots. Regional legal roots can differ while identifiers, credentials, presentations, wallet interfaces, status semantics, policy inputs and cryptographic profiles remain interoperable.

9 Reference Architecture for Verifiable Agentic Networks

The combined architecture contains eight layers. Each layer answers a distinct trust question and supplies evidence to the policy decision and enforcement loop. Governance spans several autonomous trust domains rather than one universal hierarchy.

Normative source — joint research: The following layered architecture and transaction pattern are proposals of this article. They integrate the documented M-Trust loop [1] with the Spherity identity, authority and assurance model [25–27].

Figure 3 integrates the trust domains, wallets, VDR, verifier, evidence and enforcement components into a single B2B transaction architecture for authorised enterprise agents.

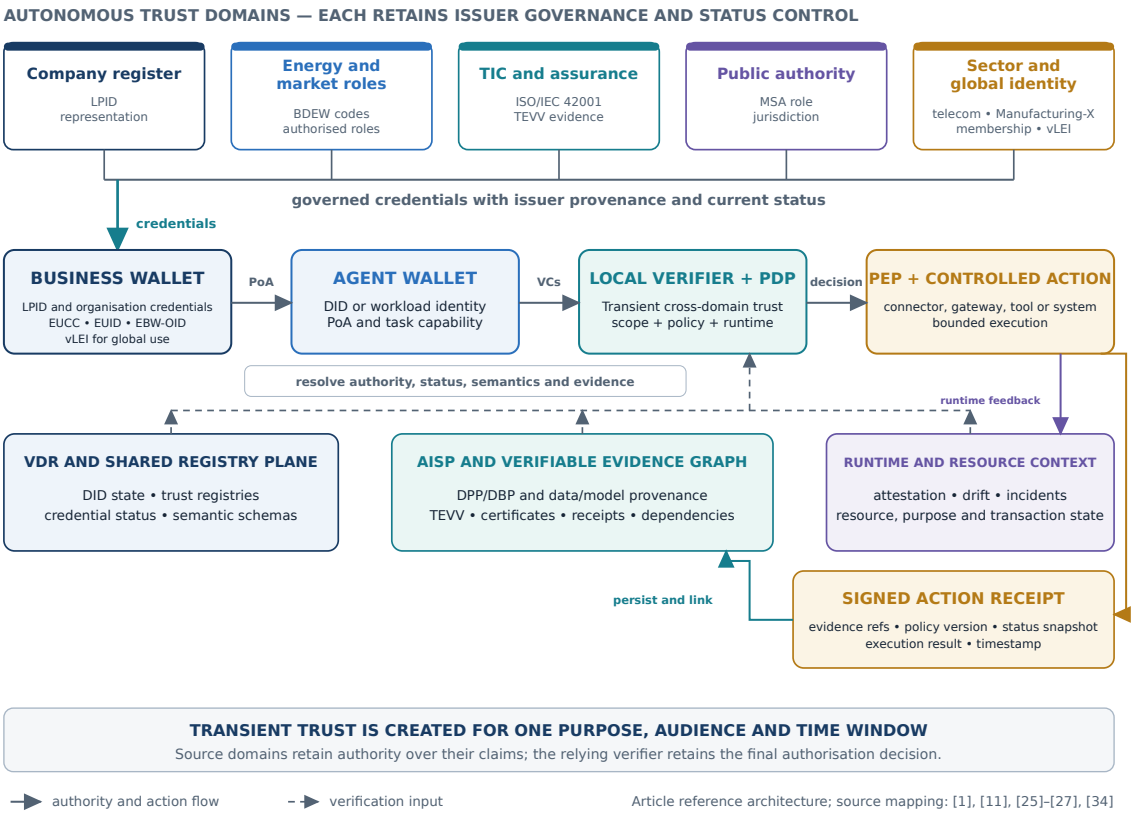

Figure 3: Reference architecture for transient cross-domain trust and verifiable agent actions for B2B use cases. Source: Author's proposal for an energy data space integrating the multi-trust loop for sector-coupling with the Spherity identity, authority and assurance model [1, 11, 25–27, 34].

9.1 Multi-Trust-Domain Governance and Transient Trust

A **trust domain** is a bounded governance environment with its own authoritative issuers, schemas, assurance rules, status processes and liability allocation. Company registers, BDEW market roles, TIC certification, government MSA roles, dena energy experts, Manufacturing-X membership, GLEIF vLEI and telecom operator identity are separate domains. A transaction can require evidence from several of them.

Cross-domain trust remains verifier-sovereign. The relying domain discovers the relevant source-domain authority through governed trust registries, verifies each credential and status entry, maps semantics into local policy and accepts the evidence for a defined purpose, audience and time. This produces **transient cross-domain trust**: a bounded decision relationship for the transaction, with each source domain retaining control of its own claims and revocation.

A VDR provides the discovery and shared-state plane. Depending on governance and legal requirements, it can be a conventional authoritative registry, permissioned DLT, EBSI-based service or qualified electronic ledger. Its functions include DID state, trust-domain and issuer entries, privacy-preserving status, semantic-scheme references and integrity or ordering commitments. It supplies verifiable inputs to Zero Trust and M-Trust; it does not replace issuer governance or the local authorisation decision.

Layer	Core components	Required outcome
1. Governance and trust	Laws, domain rulebooks, reciprocal recognition, trust registries, issuer accreditation and semantic governance	Recognised authority across autonomous trust domains
2. Organisational identity	Business Wallet, EUCC/EUID/EBW-OID, vLEI, roles, licences, signatures and seals	Verifiable legal person and representation across jurisdictions
3. Agent identity and delegation	Agent Wallet, DID/workload identity, PoA and task capability	Bounded technical agency
4. AI service assurance	AISP, TEVV, risk profile, incidents and monitoring	Verifiable system provenance and current assurance
5. Data and product evidence	Data spaces, DPP and DBP systems, provenance, semantic validation and evidence graph	Verifiable basis for decisions
6. Decision and enforcement	M-Trust evaluation, policy decision point (PDP), policy enforcement point (PEP), connector or gateway	Local, context-aware and transient cross-domain authorisation
7. Registry and cryptographic resilience	VDR, DLT/QEL/EBSI profile, hardware security modules, trusted execution environments, remote attestation, crypto agility and a PQC Corridor	Current shared state and durable integrity, authenticity and confidentiality
8. Accountability evidence	Signed action receipt, timestamp, status snapshot and audit trail	Reconstructable decision and execution history

9.2 End-to-End Authorisation Transaction

The core transaction pattern is:

1. The agent proposes a specific action with declared purpose and parameters.

2. The gateway requests the required authority and assurance evidence.

3. The Agent Wallet presents identity, PoA and relevant credentials.

4. The verifier resolves issuers, DIDs, cross-domain trust-registry entries, status services, semantic schemes and evidence references through the applicable VDRs.

5. The policy engine establishes transient trust for the required source-domain evidence and evaluates effective authority, M-Trust risk signals, AISP/TEVV state, data-space rules and runtime context.

6. The enforcement point executes the permitted operation within its bounded capability.

7. The gateway issues a signed action receipt and updates continuous monitoring.

10 Discussion: Geopolitics, Standardisation and Industrial Adoption

10.1 Source Composition and Analytical Boundary

The drafting organisations named by the M-Trust paper are the AI Research Institute of CAICT, GSMA Greater China, Huawei, China Telecom, China Unicom Research Institute, Beijing University of Posts and Telecommunications and the State Information Center. All named drafting organisations are based in China or belong to GSMA's Greater China regional structure. The contributor base combines public research, national information infrastructure, telecom

operators, a major equipment vendor, academia and a GSMA regional organisation.

M-Trust describes itself as a reference architecture and research direction and expressly disclaims status as a standard specification or mandatory implementation requirement [1, p. 29]. Its source composition and publication channel support analysis of agenda-setting effects. They provide no independent proof of coordinated state direction.

10.2 Vocabulary Power

Early definitions shape later standards, procurement criteria and regulatory questions. M-Trust frames agentic-network trust around multi-party evaluation, intent-aware policy, cross-domain propagation, heterogeneous trust roots, subscriber identity modules (SIMs), embedded SIMs (eSIMs), DLT, trusted execution environments and PQC. Adoption of the vocabulary can establish the reference problem before implementation profiles are fixed.

10.3 Distribution Power

Publication in English through GSMA places a China-developed proposal directly in global mobile-industry discourse. The format addresses operators, chief technology officers, regulators and standards participants. The separate GSMA *Mobile AI* paper reinforces the standardisation context: it calls for globally unified technical standards and coordination through international standards organisations and industry alliances [3, p. 32]. The strategic effect is agenda setting, independently of any unverified coordination intent.

10.4 Infrastructure Power and China Standards 2035

Trust architecture determines who can issue identities, operate wallets, score agents, define accepted algorithms, run registries and enforce cross-domain policy. These are market-structure decisions as well as technical choices. Bruegel describes China's AI strategy as combining domestic control and technological autonomy with global diffusion, open models and international standard setting [2]. M-Trust is consistent with that pattern: a conceptual vocabulary developed by China-based public, industrial and academic actors enters a global industry forum while competing agentic-authority profiles are still developing. This is an inference about the effect of early agenda setting, rather than a claim about unpublished intent.

The M-Trust draft also sits within a longer Chinese standardisation trajectory. Spherity's 2021 analysis of **China Standards 2035**, based on strategy material released in 2020, identified a two-pronged approach: participation in international standards bodies such as ISO and ITU, combined with Chinese-led standardisation in areas including 5G, IoT, artificial intelligence, e-commerce traceability, smart manufacturing and digital identity [52]. M-Trust is consistent with this earlier interpretation. China-based actors define a trust model for 6G and agentic networks, publish it in English through GSMA and connect it to DIDs, heterogeneous trust roots, DLT and PQC. Standards thereby function as technical coordination mechanisms, market infrastructure and instruments of geopolitical influence.

10.5 A Multipolar Interoperability Response

Normative source — joint research: The following response is the joint-research multipolar interoperability proposal.

A multipolar response SHOULD combine W3C DIDs and VCs, regional authoritative trust roots, reciprocal governance, verifier sovereignty and portable evidence so that interoperability preserves a balanced distribution of control.

European actors SHOULD bring legally grounded organisational identity, verifier sovereignty, data protection, Business Wallets, qualified trust services, governed credentials, industrial safety and data-space experience into the same standards venues. North American strengths in Zero Trust, workload identity, cloud security and TEVV and Asia-Pacific strengths in telecom scale, manufacturing and digital infrastructure complete the target architecture.

Open standards create the point of convergence. W3C DIDs and VCs, open presentation protocols, explicit issuer governance, portable status, machine-readable policy and testable PQC profiles allow regional trust roots to interoperate while preserving local law and local authorisation.

10.6 Strategic Consequences of Missing Trust Infrastructure

Agent capability alone creates no production right to act. Regulated industries, sensitive business processes and critical infrastructure require evidence that an identifiable legal person authorised the action, the agent remained within mandate, the service met applicable conformity and safety conditions, the runtime preserved cyber-resilience and the decision can be reconstructed. Manual document exchange and bilateral integration cannot provide this control at agent speed or ecosystem scale.

The absence of a common trust infrastructure therefore creates a compound adoption constraint. Legal and compliance teams cannot automate controls with sufficient evidentiary quality; functional-safety owners cannot connect an autonomous decision to current assessment and operating limits; cybersecurity teams cannot revoke authority or contain compromised agents consistently across domains; and ecosystem participants cannot reuse assurance evidence. Organisations then confine AI to advisory or low-impact functions while competitors with deployable trust controls advance into production automation.

For businesses and governments, the result is a strategic late-mover risk. Lost time compounds through delayed process learning, fragmented supplier requirements, duplicated conformance costs and dependence on external identity or agent platforms. Weak infrastructure also increases security exposure when adoption proceeds through ad hoc credentials and unverifiable delegation. Trusted-AI infrastructure is therefore both a compliance control and an industrial-capability asset.

Normative source — joint research: Regulated-sector and critical-infrastructure programmes SHOULD establish verifiable legal-person identity, delegated authority, cross-domain trust registries, functional-safety and conformity evidence, cyber-resilient enforcement and action receipts before granting AI agents production authority. Public investment programmes SHOULD treat these components as reusable ecosystem infrastructure rather than project-specific integration.

11 Research and Implementation Priorities

Normative source — joint research: The following agenda is proposed by this article. It translates the source comparison into standards and implementation work.

11.1 Semantic and Governance Profiles

Regulators, standards bodies, network operators, data-space organisations and wallet providers SHOULD prioritise:

1. Define interoperable credential profiles for legal-person identity, company roles, agent PoA, AISP, runtime attestation and action receipts.

2. Define issuer accreditation and trust-list semantics for each claim type and jurisdiction.

3. Separate risk scores, reputation and behavioural trust from legal authority in every decision model.

4. Define a verifier-sovereign cross-domain protocol that transfers evidence, status and provenance.

5. Standardise scope intersection, delegation depth, re-delegation, audience, purpose, value and expiry semantics.

6. Define reciprocal trust-domain governance and time-bounded semantic mappings for transient cross-domain trust.

11.2 System Integration and Cryptographic Resilience

Implementers SHOULD:

1. Bind agent-to-agent and Model Context Protocol interactions to external authority predicates evaluated at a gateway or connector.

2. Connect Business Wallets to data-space connectors, DPP and DBP systems and sector policy engines.

3. Define AISP and TEVV profiles that support versioning, drift, incident and runtime status.

4. Establish PQC Corridors across Europe, North America and Asia-Pacific with common conformance tests.

5. Profile DLT for minimal shared state, privacy, governance, finality, recovery and PQC transition.

6. Test EBSI/VDR and qualified-ledger service profiles for DID state, trust registries, privacy-preserving status and semantic-scheme integrity.

7. Connect IPCEI-AI and Manufacturing-X reference components to a shared control-plane and verifiable-evidence-graph profile.

11.3 Validation and Empirical Research

Research programmes SHOULD:

1. Create conformance suites for issuers, wallets, verifiers, DID resolvers, status services, policy engines and action receipts.

2. Measure decision latency, false-permit and false-deny rates, revocation propagation, resolver and status-service failure, delegation depth and PQC overhead.

3. Test downgrade, replay, issuer compromise, wallet compromise, stale mandate, poisoned evidence and cross-domain policy-conflict scenarios.

4. Compare the legal effect of machine-readable delegation and action receipts across jurisdictions

and sectors.

5. Run cross-border pilots in pharmaceuticals, energy, manufacturing, mobility, finance, telecoms and critical infrastructure.

12 Conclusion

M-Trust provides a valuable extension of Zero Trust for environments in which autonomous agents coordinate across networks and organisations. Its strongest contribution is the combination of multi-party evaluation, intent-aware authorisation, cross-domain trust and a continuous collection–decision–enforcement loop. The reviewed drafts explicitly name W3C DIDs, DLT and PQC. They use generic credential language while leaving W3C VC 2.0, digital wallets, Business Wallets and legally grounded agent delegation unspecified.

The telecom examples show how M-Trust can be implemented: dual human–agent identity, task-scoped resource authority, dynamic permission reclamation, SIM/eSIM and heterogeneous roots of trust, cross-operator workflows and bounded ledger functions. The need for Business Wallets and verifiable authorisation chains is a finding of this joint research rather than a stated GSMA requirement. A technical agent identity establishes which workload controls a key. A Business Wallet and governed credentials establish which legal person stands behind that workload, which representative delegated authority, which task and limits apply, and whether the authority remains current. Local verifier policy then computes the permitted action from the intersection of credential scopes, context and runtime state. The signed action receipt preserves the resulting accountability evidence.

European data spaces and Manufacturing-X demonstrate that the architecture generalises to industry. A legal-person root can be reused while credentials from company-register, energy-market, conformity-assessment, regulatory and professional trust domains remain under their own governance. Transient cross-domain trust lets a verifier compose those credentials for one DPP, supply-chain, sector-coupling or AI transaction. EUCC, EUID and EBW-OID provide European company-law and wallet anchors; vLEI provides complementary global entity and role interoperability. The Energy Data-X, Bundesanzeiger Verlag and GLEIF alignment paths illustrate how this architecture can move from concept towards governed implementation.

The PQC analysis reaches an equally bounded result. M-Trust identifies the quantum threat, lattice-based cryptography, lightweight implementation and hybrid migration while selecting no NIST or Chinese PQC standard. The United States has finalised NIST algorithms and federal profiles; Europe and Germany are coordinating migration and implementation assurance; China is operating an indigenous next-generation commercial-cryptography selection programme. A multi-polar implementation must extend crypto-agility across issuers, wallets, verifiers, DID methods, status services, VDRs, channels and archived evidence, with explicit cross-profile conformance testing.

DLT retains a bounded, optional role. EBSI supplies a relevant European shared trust-layer and node-operated VDR pattern; a qualified electronic ledger adds a regulated eIDAS trust-service status when its provider and service meet the statutory conditions. Local replication can reduce dependency on a live third-party lookup while preserving the need for issuer, governance, software and consensus assurance. DID state, cross-domain trust registries, privacy-preserving status and semantic-scheme integrity are the principal shared-state functions.

The comparison across M-Trust, Anthropic, WE BUILD and Manufacturing-X shows convergence at the level of architecture patterns rather than regional standards. M-Trust emphasises network trust and cross-domain evaluation. Anthropic emphasises enterprise agent security. WE BUILD

and Spherity emphasise legal-person identity, wallets and mandates. Manufacturing-X contributes industrial governance and connector ecosystems. Their combination produces a coherent target architecture for Trusted AI: verifiable authority, verifiable evidence, policy enforcement and bounded autonomy.

IPCEI-AI can operationalise this convergence by coupling the control plane to a verifiable evidence graph in the data plane. That graph makes provenance, conformity, TEVV, runtime state and action evidence traversable from the industrial decision back to its authoritative sources. Speed matters: organisations and governments that lack reusable trust infrastructure face delayed production adoption in regulated industry and critical infrastructure, duplicated compliance effort and dependence on external control layers.

A further research priority concerns dynamic agent systems that create short-lived agents for specific tasks. In these systems, the primary Agent System Identity should represent the governed system and bind it to an accountable legal person, operating policy and assurance profile. At the micro level, each dynamically created agent should receive a cryptographically verifiable identity and a task-bounded capability whose scope remains within the system-level mandate. A meso-level control plane should govern agent creation, delegation, interaction and termination. It should continuously process signed policy, runtime and TEVV evidence through machine-speed validation. Deviations should trigger proportionate enforcement: blocking or quarantining the affected micro-agent, revoking its capability, constraining the workflow, or suspending the Agent System authorisation when systemic non-conformance emerges. Further research is required on identity lineage for ephemeral agents, delegation depth, evidence latency, aggregate and emergent behaviour, cascading revocation, recovery and privacy-preserving supervision. The central design challenge is to preserve stable legal accountability at the macro level while authority is dynamically created, narrowed, monitored and withdrawn at the micro-agent level.

The geopolitical implication follows from the same standards gap. Early vocabulary can shape later regulation, procurement and infrastructure choices. A multipolar response requires interoperable W3C credential and identifier profiles, regionally legitimate trust roots, verifier sovereignty, explicit cryptographic profiles and shared conformance tests. These elements allow global agentic networks to interoperate while preserving local law, accountable organisational authority and a balanced distribution of control.

Declarations

Competing interests: The author is the founder and chief executive of Spherity GmbH and a co-author of source [6]. The manuscript evaluates Spherity research and architecture materials alongside external sources.

Data and materials availability: The public source corpus is cited below and linked in the accompanying Markdown edition. Source [27] is an internal architecture deck supplied for this analysis.

References

[1] GSMA Greater China (2026), Trust Paradigm Evolution for Agentic Networks — Draft for Comments, English- and Chinese-language drafts reviewed, 12 September 2026. Available at: https://www.gsma.com/about-us/regions/greater-china/gsma_resources/trust-paradigm-evolution-for-agentic-networks%EF%BC%88draft-for-comments%EF%BC%89/.

[2] García-Herrero, A.; Storella, T. (2026), China's artificial intelligence goals and strategic choices for Europe, Bruegel, 10 September 2026. Available at: https://www.bruegel.org/analysis/chinas-artificial-intelligence-goals-and-strategic-choices-europe.

[3] GSMA Greater China (2026), The Value of Mobile AI, Chinese-language white paper. Available at: https://www.gsma.com/solutions-and-impact/technologies/networks/gsma_resources/the-value-of-mobile-ai-whitepaper/.

[4] Data Spaces Support Centre (2025–2026), DSSC Blueprint. Available at: https://blueprint.dssc.eu/.

[5] Anthropic (2026), Zero Trust for AI Agents: A Security Framework for Deploying Autonomous AI Agents in the Enterprise, 27 May 2026. Available at: https://claude.com/blog/zero-trust-for-ai-agents.

[6] Magård, D.; Busch, P.; Hannemann, D.; Scalongne, W.; Fjelkner, B.; Parikh, S.; Stöcker, C.; Bailly, L. (2026), Trusted Identities for AI Agents: An Opportunity for Europe, WE BUILD non-paper, 27 February 2026. Available at: https://www.webuildconsortium.eu/trusted-identities-for-ai-agents-an-opportunity-for-europe.

[7] European Commission (2025), Proposal for a Regulation establishing European Business Wallets, COM(2025) 838 final, 19 November 2025. Available at: https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:52025PC0838.

[8] Plattform Industrie 4.0 (2026), Manufacturing-X Framework, Federal Ministry for Economic Affairs and Energy. Available at: https://www.plattform-i40.de/IP/Redaktion/EN/Standardartikel/ManufacturingX_Framework.html.

[9] Catena-X Automotive Network (2026), Expert Groups and Committees: SSI/Digital Identity, Dataspace Connectivity and Trusted AI, accessed 15 September 2026. Available at: https://catena-x.net/association/expert-groups-committees/.

[10] W3C (2022), Decentralized Identifiers (DIDs) v1.0. Available at: https://www.w3.org/TR/did-core/.

[11] W3C (2025), Verifiable Credentials Data Model v2.0. Available at: https://www.w3.org/TR/vc-data-model-2.0/.

[12] Gößling, J.-N.; Hoffmann, R.; Kießling, A.; Rohrbach, D.; Stöcker, C. (2026), "The energy data-X data ecosystem: Cross-sector and BESS data exchange enabled by the European Business Wallet and market role credentials," Energiewirtschaftliche Tagesfragen, 76(7–8), 19–22; extended analysis. Available at: https://spherity.github.io/spherity-research/energy-data-x-ebw-market-role-credentials-dpp-access-control.html.

[13] NIST (2024), FIPS 203: Module-Lattice-Based Key-Encapsulation Mechanism Standard. Available at: https://csrc.nist.gov/pubs/fips/203/final.

[14] NIST (2024), FIPS 204: Module-Lattice-Based Digital Signature Standard. Available at: https://csrc.nist.gov/pubs/fips/204/final.

[15] NIST (2024), FIPS 205: Stateless Hash-Based Digital Signature Standard. Available at: https://csrc.nist.gov/pubs/fips/205/final.

[16] European Commission (2025), A Coordinated Implementation Roadmap for the Transition to Post-Quantum Cryptography, 23 June 2025. Available at: https://digital-strategy.ec.europa.eu/en/library/coordinated-implementation-roadmap-transition-post-quantum-cryptography.

[17] Federal Office for Information Security (BSI) (2026), Technical Guideline BSI TR-02102-1: Cryptographic Mechanisms—Recommendations and Key Lengths, version 2026-01. Available at: https://www.bsi.bund.de/SharedDocs/Downloads/EN/BSI/Publications/TechGuidelines/TG02102/BSI-TR-02102-1.html.

[18] National Security Agency (2025–2026), Post-Quantum Cybersecurity Resources and CNSA 2.0 Guidance, accessed 15 September 2026. Available at: https://www.nsa.gov/Cybersecurity/Post-Quantum-Cybersecurity-Resources/.

[19] Institute of Commercial Cryptography Standards (2025–2026), Next-Generation Commercial Cryptographic Algorithms programme and implementation self-assessment notice, accessed 15 September 2026. Available at: https://niccs.org.cn/niccs/Notice/pc/list.html.

[20] European Commission (2026), European Blockchain Services Infrastructure (EBSI): A Shared Trust Layer, accessed 15 September 2026. Available at: https://ebsi.eu/.

[21] European Parliament and Council (2024), Regulation (EU) 2024/1183 amending Regulation (EU) No 910/2014 as regards the European Digital Identity Framework, especially Articles 3(53)–(54), 45k and 45l. Available at: https://eur-lex.europa.eu/eli/reg/2024/1183/oj/eng.

[22] Bradner, S. (1997), RFC 2119: Key words for use in RFCs to Indicate Requirement Levels, BCP 14. Available at: https://www.rfc-editor.org/rfc/rfc2119.

[23] Leiba, B. (2017), RFC 8174: Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words, BCP 14. Available at: https://www.rfc-editor.org/rfc/rfc8174.

[24] Rose, S.; Borchert, O.; Mitchell, S.; Connelly, S. (2020), NIST SP 800-207: Zero Trust Architecture. Available at: https://csrc.nist.gov/pubs/sp/800/207/final.

[25] Stöcker, C. (2026), European Business Wallets as the Legal Control Plane for Zero Trust AI Agents, Spherity Research. Available at: https://spherity.github.io/spherity-research/ebw-zero-trust-ai-agents.html.

[26] Stöcker, C. (2026), From Trustworthy AI to Trusted, Operational Agentic Systems, Spherity Research. Available at: https://spherity.github.io/spherity-research/trusted-agentic-ai-china-eu-us-comparative-analysis.html.

[27] Spherity GmbH (undated), AI-Industry Trust Infrastructure / Trusted Industrial AI, internal architecture deck, version 13; supplied PDF modified 24 April 2026.

[28] Google Cloud (2025), 'Powering AI commerce with the new Agent Payments Protocol (AP2)', 16 September 2025. Available at: https://cloud.google.com/blog/products/ai-machine-learning/announcing-agents-to-payments-ap2-protocol.

[29] Mastercard (2026), 'When AI starts buying for you, trust becomes the product: How Verifiable Intent builds trust in agentic AI commerce', 5 March 2026. Available at: https://www.mastercard.com/global/en/news-and-trends/stories/2026/verifiable-intent.html.

[30] Visa (2026), 'Trusted Agent Protocol', Visa Developer Center, updated 7 August 2026. Available at: https://developer.visa.com/use-cases/trusted-agent-protocol.

[31] Amazon Staff (2025), 'Amazon's new Buy for Me feature helps customers find and buy products from other brands' sites', About Amazon, 3 April 2025. Available at: https://www.aboutamazon.com/news/retail/amazon-shopping-app-buy-for-me-brands.

[32] Apple Inc. (2026), 'App Intents', Apple Developer Documentation, accessed 16 September 2026. Available at: https://developer.apple.com/documentation/appintents/app-intents.

[33] Apple Inc. (2026), 'In-App Purchase', StoreKit, Apple Developer Documentation, accessed 16 September 2026. Available at: https://developer.apple.com/documentation/storekit/in-app-purchase.

[34] Stöcker, C. (2026), Evidence Graphs for Industrial AI, Spherity Research. Available at: https://spherity.github.io/spherity-research/evidence-graphs-industrial-ai-data-plane.html.

[35] United Nations Commission on International Trade Law (UNCITRAL) (2024), Model Law on Automated Contracting, United Nations. Available at: https://uncitral.un.org/en/mlac.

[36] Energie Codes und Services GmbH (2026), BDEW Codes and Code-validation web service, accessed 15 September 2026. Available at: https://bdew-codes.de/.

[37] German Energy Agency (dena) (2026), Energy Efficiency Expert List for Federal Funding Programmes, accessed 15 September 2026. Available at: https://www.energie-effizienz-experten.de/.

[38] Bundesanzeiger Verlag (2026), German Company Register, accessed 15 September 2026. Available at: https://www.unternehmensregister.de/en.

[39] International Organization for Standardization (2023), ISO/IEC 42001:2023—Information technology—Artificial intelligence—Management system. Available at: https://www.iso.org/standard/42001.

[40] Council of the European Union (2026), European business wallets: Council adopts negotiating position, 9 June 2026. Available at: https://www.consilium.europa.eu/en/press/press-releases/2026/06/09/european-business-wallets-council-adopts-negotiating-position/.

[41] European Economic and Social Committee (2026), European business wallets — INT/1110, 18 March 2026. Available at: https://www.eesc.europa.eu/en/our-work/opinions-information-reports/opinions/european-business-wallets.

[42] Couzens, B.; Stöcker, C.; Vasiliu-Feltes, I. (2026), Quantum-Resilient Organizational Identity: Governance, Business Wallets, and PQC Corridors, Spherity Research. Available at: https://spherity.github.io/spherity-research/quantum-resilient-organizational-identity.html.

[43] Stöcker, C. (2026), Securing Digital Identity and Verifiable Credential Wallets against Quantum Vulnerabilities, Spherity Research. Available at: https://spherity.github.io/spherity-research/Securing-Digital-Identity-Quantum-Vulnerabilities.html.

[44] NIST (2024), NIST Releases First 3 Finalized Post-Quantum Encryption Standards, 13 August 2024. Available at: https://www.nist.gov/news-events/news/2024/08/nist-releases-first-3-finalized-post-quantum-encryption-standards.

[45] Decentralized Identity Foundation (2024), DIDComm Messaging Specification v2.1. Available at: https://identity.foundation/didcomm-messaging/spec/v2.1/.

[46] Eclipse Dataspace Working Group (2025), Dataspace Protocol 2025-1. Available at: https://eclipse-dataspace-protocol-base.github.io/DataspaceProtocol/2025-1/.

[47] Eclipse Dataspace Working Group (2025), Decentralized Claims Protocol v1.0, stable release, 12 November 2025. Available at: https://eclipse-dataspace-dcp.github.io/decentralized-claims-protocol/.

[48] German Federal Ministry for Economic Affairs and Energy (2026), IPCEI Artificial Intelligence, accessed 15 September 2026. Available at: https://www.bundeswirtschaftsministerium.de/Redaktion/EN/Dossier/ipcei-ai.html.

[49] European Commission (2026), Important Projects of Common European Interest (IPCEI), accessed 15 September 2026. Available at: https://competition-policy.ec.europa.eu/state-aid/ipcei_en.

[50] GLEIF and contributing authors (2026), Agentic AI in Payments: Establishing Interoperable Trust and Control, Working Paper Series, version 1.0, September 2026. Available at: https://www.gleif.org/organizational-identity/research-publications/2026-08-13_agentic_ai_in_payments_v1.0-1.pdf.

[51] Wieck, O. (2026), European Business Wallet: More Than a Digital Wallet for Companies, GLEIF Blog, 17 August 2026. Available at: https://www.gleif.org/en/newsroom/blog/european-business-wallet-more-than-a-digital-wallet-for-companies.

[52] Stöcker, C. (2021), Identity in a Multi-polar World and China's 2035 Standards: On the role of Digital Identity Standards in E-Commerce, Supply Chain, Global Trade and Customs, Spherity, 6 January 2021. Available at: https://medium.com/spherity/identity-in-a-multi-polar-world-and-chinas-2035-standards-66f14865b800.

This article provides an architecture and policy analysis. Deployment requires applicable legal, cybersecurity, privacy, safety and assurance review.