

HUMAN-AI PERSON STACK

Project - Framework v0.1

Identity, Memory, Portability, Agency, Authority, Provenance and
Continuity for Personal AI

STATUS

Public Discussion Draft • Exploratory conceptual framework • Not a technical standard

Originated and authored by Joe Primus Kayau

Independent project initiated in 2026

Canonical framework home: PersonTwin.com

Version date: 1 October 2026

© 2026 Joe Primus Kayau. All rights reserved. Human-AI Person Stack Framework v0.1 is an independent conceptual work. References to third-party standards, technologies, research, organizations, product names and trademarks are for identification, commentary and comparison; they do not imply ownership, endorsement or affiliation.

Publication Note

The Human-AI Person Stack Project is an independent exploratory project initiated in 2026 to develop a clear vocabulary and layered architecture for discussing how persistent personal artificial intelligence may represent, remember, move with, act for, remain accountable to, and outlive the individuals it serves. It is intentionally technology-neutral and does not prescribe a particular model provider, cryptographic stack, payment network, government identity system or legal regime.

This v0.1 document proposes seven distinct but related concepts. It does not claim that these terms are established industry standards. The purpose of publication is to make the concepts explicit, testable, referenceable and open to critique while the surrounding AI-agent ecosystem is still forming. The seven-layer arrangement and the canonical definitions are an original synthesis written for this framework. External research, standards and commercial work are cited as background and points of comparison; titles and established terms are used only to identify those sources. This publication is not a substitute for formal copyright, trademark or patent clearance.

Important boundary

PersonPassport is not a travel passport or government credential; HumanitySeal is not a guarantee that content is true; AfterTwin is not a claim that a digital model is the deceased person. These names are conceptual labels for future system functions and governance questions.

About the Project and Originator

The Human-AI Person Stack Project is an independent conceptual project initiated in 2026. Its purpose is to organize questions about persistent personal AI into a coherent, referenceable framework for discussion, comparison, research and critique. It does not claim institutional or standards-body authority.

Joe Primus Kayau is a professional engineer and entrepreneur with an independent interest in the long-term relationship between people, artificial intelligence and digital identity. He is the originator and author of Framework v0.1. The framework is an independent cross-disciplinary synthesis and is not presented as expert legal, regulatory, standards or technical guidance. Specialist review and critique are encouraged as the framework develops.

Suggested citation: *Joe Primus Kayau (2026). Human-AI Person Stack Project: Framework v0.1. Public Discussion Draft. PersonTwin.com.*

Contents

1. Executive Summary
2. Why This Framework, Why Now
3. Design Principles
4. Architecture Overview
5. Layer 1 — PersonTwin
6. Layer 2 — PersonPassport
7. Layer 3 — PortableTwin
8. Layer 4 — PersonProxy
9. Layer 5 — TwinAuthority
10. Layer 6 — HumanitySeal
11. Layer 7 — AfterTwin
12. How the Layers Work Together
13. 2030 End-to-End Scenario
14. Sector Use Cases
15. Security, Privacy and Ethical Risks
16. Relationship to Existing Standards and Industry Work
17. Governance and Versioning
18. Publication and Adoption Strategy
19. 90-Day Roadmap
20. Open Questions for v0.2
21. Canonical Definitions
22. References

1. Executive Summary

Personal AI is moving beyond isolated chat sessions toward systems that can retain person-specific context, participate in communication, carry out delegated tasks and interact with outside services. Work published in 2025-2026 spans portable personal context, interoperable agents, verifiable credentials, agent identity and authorization, provenance systems and agent-initiated commerce.[1][2][3][4][8][9]

The Human-AI Person Stack proposes seven conceptual layers for that emerging environment. The framework separates questions that are often collapsed into the single word “agent”: representation, identity, portability, delegated action, authority, provenance and lifecycle continuity. The separation matters because a system may know a person without being authorized to act, may be authorized to act without carrying all of the person’s memories, and may prove human approval without proving that every output was written manually by a human.

Core thesis

As personal AI becomes persistent and capable of action, the durable infrastructure will not be only the intelligence model. It will also be the human-specific context, identity, permissions, provenance and continuity that survive across models and providers.

Layer	Concept	Primary question	Canonical domain
1	PersonTwin	Who is the persistent digital counterpart?	persontwin.com
2	PersonPassport	Who does it represent and what can it prove?	personpassport.com
3	PortableTwin	How does the person-specific context move?	portabletwin.com
4	PersonProxy	How does AI act on the person’s behalf?	personproxy.com
5	TwinAuthority	What is the AI allowed to do?	twinauthority.com
6	HumanitySeal	Where is human origin, review or authorization?	humanityseal.com
7	AfterTwin	What happens after death, incapacity or retirement?	aftertwin.com

2. Why This Framework, Why Now

The timing is important because technical capability is moving quickly while the language for describing responsibility and control is still unsettled. Stanford research published in 2026 examines ways for individuals to carry selected long-term context between AI tools under user control.[1] NIST work on AI-agent standards addresses interoperability and security questions, including how software agents can be identified, authenticated and authorized.[2][3]

Microsoft has also published work on AI systems designed to represent or assist specific people in communication and workplace settings.[6][7] In payments, Visa and Mastercard are developing infrastructure that allows software agents to participate in commerce within user-set constraints, identity checks and consent mechanisms.[8][9][10]

These developments suggest that by 2030 the central challenge may not be whether AI can perform useful tasks, but whether people and institutions can answer: Who is this AI representing? What context does it carry? What credentials does it have? What is it authorized to do? Can its actions be audited and revoked? What evidence exists of human approval? What happens to its accumulated human context over a lifetime?

Why publish in 2026

The framework is early enough to be exploratory, but late enough to be grounded in real research, standards and commercial activity. A dated v0.1 creates a reference point without pretending that terminology or implementations are settled.

3. Design Principles

Human sovereignty: The individual remains the primary source of authority over personal context, delegation and lifecycle choices.

Least authority: An AI should receive only the permissions needed for a defined task, for a defined duration and under defined conditions.

Portability without overexposure: Moving context should not mean copying an entire life history into every service.

Separation of identity and delegated power: Establishing whom a system represents is a different question from deciding what that system may do.

Revocability: Delegated access and authority should be withdrawable, including during compromise, incapacity or changing circumstances.

Auditability: Consequential actions should be attributable to the system, policy and human authorization chain that produced them.

Human legibility: Permissions and provenance must be understandable to ordinary people, not only to developers.

Interoperability: Concepts should map to open standards where practical rather than recreate credentials, provenance or authorization mechanisms from scratch.

Lifecycle continuity: Systems should account for changes across years: new providers, new devices, changing capabilities, incapacity and death.

Non-deception: A digital representation should not be confused with the human person it models or represents.

4. Architecture Overview

The stack is conceptual rather than implementation-specific. A real product might combine several layers in one application, or use multiple vendors. The framework separates the functions so that identity, memory, authority and provenance can be reasoned about independently.

Human-AI Person Stack

Seven conceptual layers for persistent personal AI




Figure 1. Human-AI Person Stack v0.1. The arrows represent a conceptual lifecycle and dependency flow, not a mandatory technical sequence.

Key distinction

A PersonTwin may exist without a PersonProxy. A PersonProxy should not act without TwinAuthority. A PersonPassport may establish identity without exposing full memory. HumanitySeal may prove human approval without proving the underlying content is factually correct.

5. Layer 1 — PersonTwin

Canonical definition

A persistent, person-specific AI context and representation that accumulates corrected memory, preferences, relationships, knowledge and goals across time while remaining distinguishable from the human it represents.

Primary question: Who is the digital counterpart?

Purpose and scope

PersonTwin is the anchor layer. It describes continuity of person-specific context across many model versions rather than a single chatbot session. A PersonTwin may include structured preferences, memory, communication style, trusted relationships, professional context, goals and historical decisions. The concept is compatible with many implementations: a local personal model, a cloud agent, a federated memory layer, or an ensemble of services.

Boundary

A PersonTwin is not necessarily autonomous. It may know the person deeply yet have no authority to transact, sign, disclose sensitive information or speak externally. Those functions belong to other layers.

Why the concept is timely

Microsoft has published examples of AI systems intended to represent or assist specific people in mediated communication and work contexts.[6][7]

Illustrative capabilities

- Persistent preference memory
- Relationship and role context
- Long-term goals and constraints
- Personal communication style
- Versioned memory and corrections

Canonical domain: persontwin.com

6. Layer 2 — PersonPassport

Canonical definition

A portable trust package that lets a personal AI prove whom it represents and present only the credentials or attributes needed for a particular interaction.

Primary question: Who does it represent and what can it prove?

Purpose and scope

PersonPassport is the trust and identity layer. It may reference verified credentials, account bindings, device attestations, role memberships or other claims that can be selectively disclosed. Its purpose is not to reproduce a government passport, but to provide a conceptual container for identity and credentials in the human-AI relationship.

Boundary

PersonPassport should not automatically contain the full PersonTwin memory store. Identity proof and personal history are different risk domains and should be separable.

Why the concept is timely

W3C Verifiable Credentials 2.0 provides a standards-based way to express and verify digital credentials. PersonPassport can build conceptually on such machinery rather than inventing a parallel credential format.[4][5]

Illustrative capabilities

- Selective credential disclosure
- Verified account/person bindings
- Role and entitlement proofs
- Portable trust metadata
- Credential revocation and expiry

Canonical domain: personpassport.com

7. Layer 3 — PortableTwin

Canonical definition

A controlled portability capability that lets selected parts of a person's persistent AI context move or synchronize across providers, devices and intelligent environments without requiring the entire context store to be copied.

Primary question: How does the person-specific context move?

Purpose and scope

PortableTwin is the interoperability layer. It addresses vendor lock-in, continuity across model changes, migration between services and selective sharing of personal context. Portability should support role- and purpose-specific subsets rather than unrestricted copying of a person's entire history.

Boundary

Portability is not identical to identity. A system may verify the person through PersonPassport yet receive only a small, purpose-specific slice of context.

Why the concept is timely

Stanford research on portable personal memory and the Human Context Protocol examines how selected long-term context could be carried or made available across tools under user control.[1][11]

Illustrative capabilities

- Export/import of personal context
- Purpose-specific context bundles
- Cross-provider migration
- Interoperability profiles
- User-controlled deletion and synchronization

Canonical domain: portabletwin.com

8. Layer 4 — PersonProxy

Canonical definition

A personal AI delegate designated to perform a bounded set of actions for an identified individual, with scope, duration, limits and revocation conditions that can be inspected.

Primary question: How does AI act on the person's behalf?

Purpose and scope

PersonProxy is the agency layer. Examples include communicating, arranging travel, requesting quotes, managing subscriptions, purchasing within limits, coordinating appointments, or interacting with other agents. The defining property is representation: the system acts for a particular human principal.

Boundary

PersonProxy is not unrestricted autonomy. Delegation should be task-specific, bounded and observable, with escalation for consequential actions.

Why the concept is timely

Visa and Mastercard have announced agentic-payment work involving agent identification, user consent, transaction controls and delegated purchasing.[8][9][10]

Illustrative capabilities

- Travel and appointment booking
- Shopping and negotiation
- Routine communications
- Subscription and household administration
- Agent-to-agent coordination

Canonical domain: personproxy.com

9. Layer 5 — TwinAuthority

Canonical definition

The permission framework that states which actions a personal AI may undertake, on whose behalf, under what conditions, for how long, and with what approval or audit requirements.

Primary question: What is the AI allowed to do?

Purpose and scope

TwinAuthority separates intelligence from power. A highly capable model may receive minimal authority, while a simpler agent may receive narrow operational permissions. Authority can be expressed through scopes, budgets, durations, approval thresholds, counterparties, data domains and emergency revocation rules.

Boundary

TwinAuthority should never be inferred merely from the richness of the PersonTwin. Knowing a person well is not permission to act broadly for them.

Why the concept is timely

NIST publications in 2026 address identity and authorization for software and AI agents.[2][3] Visa's agentic-commerce materials likewise describe controlled credentials and permissions for agents participating in transactions.[8]

Illustrative capabilities

- Spending limits
- Data-access scopes
- Counterparty allow/deny lists
- Human approval thresholds
- Time-bounded and revocable permissions

Canonical domain: twinauthority.com

10. Layer 6 — HumanitySeal

Canonical definition

A machine-verifiable record of the nature of human involvement in a digital artifact or action - for example, originated by, reviewed by or expressly authorized by a person.

Primary question: Where is human origin, review or authorization?

Purpose and scope

HumanitySeal is the provenance and accountability layer. It should distinguish different claims: “created by a human,” “reviewed by a human,” “approved by a human,” and “generated by AI under human authority.” Those are not equivalent.

Boundary

A HumanitySeal should not be treated as proof that content is true, wise, lawful or ethical. It is a statement about provenance or human participation, not factual correctness.

Why the concept is timely

C2PA provides a technical ecosystem for recording content provenance and authenticity information. HumanitySeal is positioned as a semantic layer for describing the nature of human involvement, not as a competing provenance protocol.[12][13]

Illustrative capabilities

- Human-origin assertion
- Human-review assertion
- Human-approval assertion
- Authorization timestamp
- Link to verifiable provenance evidence

Canonical domain: humanityseal.com

11. Layer 7 — AfterTwin

Canonical definition

The stewardship rules for personal AI representations, memories, access rights and delegated functions when the represented person dies, becomes permanently incapacitated or intentionally retires the system.

Primary question: What happens when the person can no longer govern the system?

Purpose and scope

AfterTwin is the lifecycle-continuity layer. It addresses preservation, deletion, family access, archival use, executorship, research use, memorial experiences and the termination of delegated authority. It also covers incapacity scenarios in which the person is alive but temporarily or permanently unable to manage the system directly.

Boundary

AfterTwin must not imply that a digital system is the deceased person. It is a governed representation built from traces, memories and models, subject to prior consent and post-life stewardship rules.

Why the concept is timely

This layer becomes more important as PersonTwins accumulate years or decades of context. It should ultimately connect to estate planning, digital legacy law, consent directives and technical deletion/retention mechanisms.

Illustrative capabilities

- Legacy directives
- Executor/steward roles
- Memory archive modes
- Posthumous access rules
- Shutdown, deletion or preservation choices

Canonical domain: aftertwin.com

12. How the Layers Work Together

The value of the framework comes from interaction between layers. A real human-AI system should be able to answer the following chain of questions independently:

PersonTwin — What does this AI know about the person?

PersonPassport — How can another system verify whom it represents?

PortableTwin — What context may move to a new system?

PersonProxy — What action is being performed on the person's behalf?

TwinAuthority — Is that action within the granted scope?

HumanitySeal — What evidence exists of human origin, review or approval?

AfterTwin — What happens when the human can no longer govern the system?

Design rule

No layer should silently imply another. Proof of identity is not a grant of delegated power.

Possession of memory is not consent. Human review is not a truth guarantee. Portability should be selective rather than wholesale replication. Preserving a representation does not authorize it to keep operating.

13. 2030 End-to-End Scenario

Consider a user named Amina in 2030. Her personal AI has accompanied her across work, family life, travel and healthcare for several years.

1. **PersonTwin:** Amina's persistent context contains preferences, frequent contacts, travel habits, accessibility needs and long-term goals.
2. **PersonPassport:** When Amina connects a new banking agent, the bank receives only the verified credentials needed to establish identity and eligibility.
3. **PortableTwin:** Amina changes AI providers. She exports an approved subset of her personal context rather than rebuilding preferences from scratch.
4. **PersonProxy:** She asks her proxy to arrange a family trip, compare flights and hotels, coordinate calendars and purchase within a budget.
5. **TwinAuthority:** The proxy may spend up to a defined amount with approved merchants, but any larger transaction requires explicit approval.
6. **HumanitySeal:** A final high-value purchase records that Amina personally approved the transaction, distinct from the AI-generated itinerary.
7. **AfterTwin:** Years later, Amina's estate directive specifies which memories may be archived for family, which must be deleted, and that all transactional authority terminates immediately on death.

14. Sector Use Cases

Financial services

- Identity-aware agent access
- Spending limits and approval thresholds
- Agentic commerce and recurring purchases
- Audit trails and non-repudiation
- Fraud controls for impersonating agents

Healthcare and care coordination

- Portable patient preferences
- Care-plan memory with purpose limitation
- Appointment and insurance administration
- Human approval for sensitive disclosures
- Incapacity directives and caregiver authority

Workplace

- Personal digital twin for meetings and knowledge
- Delegated scheduling and routine communication
- Role-specific enterprise credentials
- Authority boundaries for procurement or approvals
- Separation of employee memory from employer ownership

Family and ageing

- Long-term personal context
- Caregiver/delegate roles
- Safe financial administration
- Memory preservation
- Legacy and succession controls

Education

- Portable learning history
- Student-controlled learning preferences
- Credential portability
- AI tutoring with scoped memory
- Transition across institutions

Government and civic services

- Verified eligibility and identity claims
- Delegated administrative assistance
- Explicit consent for data sharing
- Accessible service navigation
- Strong separation between person identity and agent authority

15. Security, Privacy and Ethical Risks

Impersonation: An attacker may attempt to operate a fake PersonTwin or PersonProxy. Mitigation requires strong bindings between human identity, agent identity and authorization state.

Memory poisoning: Incorrect or maliciously inserted memories can distort future decisions. Systems need provenance, correction workflows and confidence metadata for memory.

Over-delegation: Users may grant broad authority for convenience. TwinAuthority should favor least privilege, clear expiry and high-friction escalation for consequential acts.

Privacy leakage: Portability can become mass replication. PortableTwin should support selective, purpose-bound context rather than indiscriminate export.

Vendor lock-in: A provider may make memory technically exportable but practically unusable elsewhere. Interoperability must include semantics, not only file export.

Anthropomorphic confusion: People may over-attribute consciousness, intent or continuity to a model. The framework requires clear distinction between human and representation.

Posthumous harm: AfterTwin systems could distress families, misrepresent the deceased, or operate contrary to prior wishes. Explicit lifetime directives and stewardship are essential.

Unequal access: High-quality personal AI could amplify socioeconomic inequality. Public-interest and open standards work should consider affordability and accessibility.

Institutional power imbalance: Employers, platforms or governments could pressure people to disclose more personal context than necessary. Selective disclosure and purpose limitation are core safeguards.

False trust signals: HumanitySeal could be misread as “true” or “safe.” Its semantics must specify what was actually verified: origin, review, approval or authorization.

16. Relationship to Existing Standards and Industry Work

The Human-AI Person Stack is best understood as a conceptual coordination layer. It should map to existing technical standards wherever possible and introduce new protocol work only where genuine gaps remain.

External work	Relevant stack concepts	Relationship
Stanford Portable Memory / Human Context Protocol	PersonTwin, PortableTwin	Related work on carrying selected long-term personal context across tools under user control.[1][11]
NIST AI Agent Standards Initiative	PersonProxy, TwinAuthority, PersonPassport	Related work on interoperable agents and the identity, authentication and authorization controls around them.[2][3]
W3C Verifiable Credentials 2.0	PersonPassport	Potential technical basis for presenting and verifying digital credentials or claims.[4][5]
C2PA Content Credentials	HumanitySeal	Potential provenance substrate; HumanitySeal would describe the meaning of human participation rather than replace C2PA.[12][13]
Microsoft Dittos / DigitalMe	PersonTwin, PersonProxy	Examples of research and workplace experimentation involving AI representations or assistants for specific people.[6][7]
Visa Intelligent Commerce	PersonProxy, TwinAuthority, PersonPassport	Commercial work on agent participation in payments with identity, permissions and user controls.[8]
Mastercard Agent Pay	PersonProxy, TwinAuthority	Commercial work on registered agents, consent signals and agent-initiated payment flows.[9][10]

17. Governance and Versioning

The framework should remain lightweight until external use justifies formal governance. v0.1 should establish stable definitions and a transparent change history without claiming standards authority.

- Each canonical concept has one public definition and one canonical domain.
- Definitions are versioned. Changes that materially alter meaning require a new minor or major version.
- Implementation guidance is separated from concept definitions so that terminology can survive changing technology.
- External critiques and proposed amendments should be logged publicly.
- No claim of conformance should be made before a future technical profile defines testable requirements.
- The framework should avoid duplicating established credential, provenance or authorization standards.
- Before any formal trademark, standards-body submission or patent-sensitive technical disclosure, obtain professional legal/IP advice.

18. Publication and Adoption Strategy

The objective of v0.1 is not to “win” terminology immediately. It is to create a coherent, dated, useful reference that others can understand and cite. Adoption should be earned through clarity and utility rather than asserted.

PersonTwin.com as framework home: Add a “Human-AI Person Stack” section that links all seven concepts and hosts this paper in HTML and PDF form.

Canonical concept pages: Expand each of the seven landing pages with the v0.1 definition, scope, boundary, examples and relationship to the other six layers.

Version history: Publish a visible changelog so future readers can distinguish the original 2026 definitions from later revisions.

Public repository: Create a repository containing definitions, diagrams, references and machine-readable metadata. Keep technical code minimal until a real implementation is justified.

Research notes: Publish occasional short analyses when major developments in personal AI, agent identity, portability, provenance or digital legacy directly affect the framework.

Feedback channel: Use the strategic@ addresses for research and partnership enquiries and add a dedicated framework feedback address later if volume grows.

Archival publication: After legal/IP review, consider a stable public archive such as Zenodo or OSF so each release has a persistent citation record.

19. 90-Day Roadmap

Period	Recommended action
Days 1–14	Publish Framework v0.1 PDF + HTML; cross-link all seven sites; add definitions and “not a standard” status labels.
Days 15–30	Create public repository; add version history; add a one-page architecture image and downloadable reference sheet.
Days 31–60	Publish 3–5 research notes mapping real developments to the stack; invite targeted feedback from identity, payments, digital-twin and AI-agent practitioners.
Days 61–90	Review feedback; perform trademark/name screening; identify terminology that needs revision; publish v0.2 only if changes materially improve clarity.
Trigger for v1.0	Do not rush. Consider v1.0 only after external references, one or more implementation experiments, and a clearer mapping to standards and legal requirements.

20. Open Questions for v0.2

- What is the minimum viable data model for a PersonPassport without turning it into a universal dossier?
- Which parts of PersonTwin memory should be portable by default, portable only by explicit choice, or non-portable?
- How should TwinAuthority express nested delegation when a PersonProxy calls other agents?
- What human-approval semantics should HumanitySeal standardize first: origin, review, approval, or authorization?
- How can users inspect and correct memory provenance before it influences consequential decisions?
- What legal form should posthumous directives take for AfterTwin systems?
- How should authority change during temporary incapacity, guardianship, divorce, employment changes or account compromise?
- What terminology will ordinary consumers naturally adopt by 2030, and where should the framework adapt rather than defend a weak term?
- Which components should be open specifications, reference implementations, or purely descriptive concepts?
- How can the framework avoid becoming a mechanism for excessive surveillance or coerced disclosure?

21. Canonical Definitions — v0.1

PersonTwin

A persistent, person-specific AI context and representation that accumulates corrected memory, preferences, relationships, knowledge and goals across time while remaining distinguishable from the human it represents.

PersonPassport

A portable trust package that lets a personal AI prove whom it represents and present only the credentials or attributes needed for a particular interaction.

PortableTwin

A controlled portability capability that lets selected parts of a person's persistent AI context move or synchronize across providers, devices and intelligent environments without requiring the entire context store to be copied.

PersonProxy

A personal AI delegate designated to perform a bounded set of actions for an identified individual, with scope, duration, limits and revocation conditions that can be inspected.

TwinAuthority

The permission framework that states which actions a personal AI may undertake, on whose behalf, under what conditions, for how long, and with what approval or audit requirements.

HumanitySeal

A machine-verifiable record of the nature of human involvement in a digital artifact or action - for example, originated by, reviewed by or expressly authorized by a person.

AfterTwin

The stewardship rules for personal AI representations, memories, access rights and delegated functions when the represented person dies, becomes permanently incapacitated or intentionally retires the system.

Canonical Domains

persontwin.com — persistent representation <https://persontwin.com>

personpassport.com — identity and credentials <https://personpassport.com>

portabletwin.com — portability and interoperability <https://portabletwin.com>

personproxy.com — delegated action <https://personproxy.com>

twinauthority.com — permissions and limits <https://twinauthority.com>

humanityseal.com — human provenance and authorization <https://humanityseal.com>

aftertwin.com — continuity and stewardship <https://aftertwin.com>

Framework provenance: Human-AI Person Stack Framework v0.1 was originated and authored by Joe Primus Kayau and first prepared for public discussion in 2026. The seven-layer arrangement and canonical definitions are presented as an independent synthesis; cited third-party work remains the property of its respective authors and organizations.

22. References

- [1] Stanford Digital Economy Lab (2026). The Future of Personal AI: Portable and Persistent Personal Memory through a Unified Human Context Protocol. <https://digitaleconomy.stanford.edu/publication/the-future-of-personal-ai-portable-and-persistent-personal-memory-through-a-unified-human-context-protocol/>
- [2] National Institute of Standards and Technology (NIST) (2026). AI Agent Standards Initiative. <https://www.nist.gov/artificial-intelligence/ai-agent-standards-initiative>
- [3] NIST National Cybersecurity Center of Excellence (2026). Accelerating the Adoption of Software and Artificial Intelligence Agent Identity and Authorization. <https://csrc.nist.gov/pubs/other/2026/02/05/accelerating-the-adoption-of-software-and-ai-agent/ipd>
- [4] World Wide Web Consortium (2025). The Verifiable Credentials 2.0 family of specifications is now a W3C Recommendation. <https://www.w3.org/news/2025/the-verifiable-credentials-2-0-family-of-specifications-is-now-a-w3c-recommendation/>
- [5] World Wide Web Consortium (2025). Verifiable Credentials Data Model v2.0. <https://www.w3.org/TR/vc-data-model/>
- [6] Microsoft Research (2026). Dittos: Mimetic, Reciprocal Agents in AI Mediated Communication. <https://www.microsoft.com/en-us/research/publication/dittos-mimetic-reciprocal-agents-in-ai-mediated-communication/>
- [7] Microsoft Inside Track (2026). Meet DigitalMe: Our AI digital twin that works on our behalf. <https://www.microsoft.com/insidetrack/blog/meet-digitalme-our-ai-digital-twin-that-works-on-our-behalf/>
- [8] Visa (2026). Agentic AI enables businesses to transact with confidence. <https://corporate.visa.com/en/solutions/intelligent-commerce/vcs-agentic-ai.html>
- [9] Mastercard (2026). Mastercard Agent Pay: secure, scalable and trusted agentic AI. <https://www.mastercard.com/us/en/business/artificial-intelligence/mastercard-agent-pay.html>
- [10] Mastercard Newsroom (2026). Mastercard conducts first live agentic transaction in Malaysia with CIMB and RHB Pilot. <https://www.mastercard.com/news/ap/en/newsroom/press-releases/en/2026/mastercard-conducts-first-live-agentic-transaction-in-malaysia-with-cimb-and-rhb-pilot/>
- [11] Stanford Digital Economy Lab (2026). Human Context Protocol (HCP). <https://digitaleconomy.stanford.edu/project/loyal-agents/hcp-human-context-protocol/>
- [12] Coalition for Content Provenance and Authenticity (2026). C2PA Specifications. <https://spec.c2pa.org/specifications/>
- [13] Coalition for Content Provenance and Authenticity (2026). Content Credentials: C2PA Technical Specification 2.4. https://spec.c2pa.org/specifications/specifications/2.4/specs/C2PA_Specification.html

Next revision

v0.2 should be driven by external critique, terminology testing and standards mapping — not by adding more concepts for their own sake. The framework should remain small enough that each layer answers a distinct question.