

VAOM Alignment Annex

Regulatory, standards, and market alignment for the Verkflöde Agent Operating Model

Companion to VAOM v5.0. Revised quarterly without a framework
version change.

// Facts verified against primary sources as of 27 September 2026.

DELEGATION-FIRST · ARCHITECTURE-BACKED · COMPLIANCE-CREDIBLE · ENFORCEMENT-READY ·
EVIDENCE-PROVEN

Annex revision 2026-10 · October 2026

Verkflöde AB · verkflode.com/vaom-annex

Licensed under Creative Commons Attribution 4.0 International (CC BY 4.0)

Contents

	About This Annex	3
A1	European Union	3
A2	United Kingdom	6
A3	United States	7
A4	Asia-Pacific and International	8
A5	Runtime Control and Agent Security Standards	10
A6	Identity, Authorization, and Agent Protocols	12
A7	Agent Payments	13
A8	Liability and Insurance	14
A9	Research Alignment	14
	Annex References	15
	Annex Revision History	19

About This Annex

The VAOM whitepaper describes a method and an architecture that are meant to stay stable across releases. The landscape it operates in does not stay still: application dates move, drafts become standards, protocols merge, and products ship. Earlier versions carried both in one document, which meant that every change in the landscape forced a new framework version, and that the parts of the paper most likely to age sat beside the parts that should not.

From version 5.0, the whitepaper's Section 13 states the regulatory design obligations VAOM is built to meet, in terms that do not depend on dates. This annex carries everything that does: the current status of regulations, standards, protocols, and products, and how each maps to VAOM. It is revised on its own schedule, quarterly as a rule, without a framework version change. Each revision is dated, and each factual claim is verified against a primary source or marked as reported where only secondary sources were available.

Nothing in this annex is legal advice. Whether a given obligation applies to a given system depends on its classification and processing context, which VAOM does not determine.

A 1

European Union

The AI Act after the Digital Omnibus

Regulation (EU) 2026/1744, the Digital Omnibus on AI, was signed on 8 July 2026, published in the Official Journal on 24 July 2026, and entered into force on 27 July 2026.¹ It changes when the AI Act's obligations apply more than what they require.²

PROVISION	STATUS AFTER THE OMNIBUS
High-risk systems under Annex III (stand-alone)	Apply from 2 December 2027
High-risk systems under Annex I (in regulated products)	Apply from 2 August 2028
Transparency: Article 50(1), (3), (4)	Apply from 2 August 2026, unchanged
Transparency marking: Article 50(2)	Systems already on the market before 2 August 2026 have until 2 December 2026
New prohibitions: Article 5(1)(ba) and (bb)	Non-consensual intimate imagery and child sexual abuse material; apply from 2 December 2026

PROVISION	STATUS AFTER THE OMNIBUS
AI literacy: Article 4	Changed from ensuring a sufficient level to taking measures to support AI literacy
Bias detection: new Article 4a	Permits exceptional processing of special category data to detect and correct bias in high-risk systems
AI Office	Exclusively competent to supervise certain AI systems: those built on a general-purpose model by the same provider (with carve-outs), and those that are or are integrated into designated very large online platforms and search engines
Legacy high-risk systems used by public authorities	Must comply by 2 August 2030 (Article 111(2), retained)
Supplier agreements: Article 25	Written agreements now cover suppliers of AI models as well as tools, services, and components; the original provider must inform new providers of known limitations and failure modes

For agents specifically, the Commission's AI Act Service Desk states that agents are not a separate category: the definitions of an AI system and of a general-purpose AI model "are sufficient to cover AI agents", and the Commission's regulatory considerations on agents are "only preliminary at this stage".³ The term appears in the Act's text for the first time in the Omnibus's new Annex XIV, where code AIH 0401 ("AI systems based on other emerging AI technologies not covered by other codes, including Agentic AI") lets conformity assessment bodies be designated for agentic systems. It is a scope code for notified bodies, not a legal category or an obligation.

Guidance is arriving unevenly. Draft guidelines on high-risk classification under Article 6 were published for targeted consultation on 19 May 2026 and do not address agents.⁴ Guidelines on the Article 50 transparency obligations were adopted on 20 July 2026,⁵ and a voluntary Code of Practice on marking and labelling AI-generated content was published on 10 June 2026.⁶ We found no Commission guidelines specific to Article 14 human oversight as of September 2026.

Mapping to VAOM. Article 50(1) matters now, not in 2027: agents that interact directly with people, such as the customer-facing flows of worked example 7A, must be designed so that people know they are dealing with an AI system unless it is obvious. The obligation sits with the provider. Article 14(4)(b), which requires that overseers be enabled to remain aware of automation bias, is the provision Readiness Condition 6 and the oversight-health metrics are built to evidence. Article 25's supplier agreements are the contractual counterpart of Section 12's vendor-model channel.

Harmonised standards

STANDARD	SUPPORTS	STATUS (SEPTEMBER 2026)
EN 18286:2026 <i>Quality management system for EU AI Act regulatory purposes</i>	Article 17	Published by CEN-CENELEC in July 2026. ⁷ Presumption of conformity arises only once its reference is published in the Official Journal, which we had not found
prEN 18228 <i>Risk management</i>	Article 9	Draft; went to CEN Enquiry in 2026 (reported) ⁸
prEN 18229-1 (logging)	Article 12	Draft; reported to have passed CEN Enquiry

STANDARD	SUPPORTS	STATUS (SEPTEMBER 2026)
prEN 18229-3 <i>AI trustworthiness framework, Part 3: Human oversight</i>	Article 14	Draft at CEN Enquiry; national comment period 3 August to 7 October 2026 ⁹

One commentator on the non-public enquiry draft of prEN 18229-3 reports that it organizes oversight measures around a "reaction timeframe", the time available for a person to intervene before harm occurs.¹⁰ If that survives into the published standard, it will be a natural input to Band B review SLAs and to the choice between Draft & Approve and Execute & Audit. Until then it is reported, not established.

GDPR Article 22 and explanation

Article 22(1) GDPR is, in the Court of Justice's words, "a prohibition in principle" on decisions based solely on automated processing that produce legal or similarly significant effects (*SCHUFA*, C-634/21, 2023)¹¹. The Court held that a probability value produced automatically is itself such a decision where a third party "draws strongly" on it. VAOM draws an inference by analogy, not a holding: an agent output that a human reviewer approves without meaningful review is at risk of being treated as the decision. The Article 29 Working Party's guidance, endorsed by the EDPB, requires that human involvement be "meaningful, rather than just a token gesture" and "carried out by someone who has the authority and competence to change the decision".¹³

Explanation rests on Article 15(1)(h). In *Dun & Bradstreet Austria* (C-203/22, 2025)¹² the Court held that the data subject may require "the procedure and principles actually applied" to be explained in a concise and intelligible form, with any trade secrets disclosed to the supervisory authority or court for balancing rather than withheld.

The law may change. The Commission's data-side Digital Omnibus proposal, COM(2025) 837 of 19 November 2025,¹⁴ would recast Article 22(1) as a permission: solely automated decisions would be allowed where necessary for a contract, "regardless of whether the decision could be taken otherwise than by solely automated means", where authorised by law, or with explicit consent. The EDPB and EDPS, in Joint Opinion 2/2026, asked that the wording keep the form of a prohibition with exceptions.¹⁶ As of September 2026 the proposal is at committee stage in the European Parliament (draft report of 22 June 2026, amendments tabled 27 July 2026), with no committee vote, no mandate, and no trilogue.¹⁵ None of its GDPR changes has been adopted. The data-side Omnibus should not be confused with the AI Omnibus, which is law; the EDPB-EDPS opinion on the AI Omnibus is Joint Opinion 1/2026.

Mapping to VAOM.

OBLIGATION	WHAT IT DEMANDS OF A DELEGATION DESIGN	VAOM MECHANISM AND EVIDENCE
Article 22(1): no solely automated decision with legal or similarly significant effect, unless an exception applies	Identify every such decision; where no exception applies, it must not be solely automated	Decision Inventory flags it; Dimension 3 classifies it as regulated or prohibited from full automation; Band A is structurally unavailable, enforced by tool or scope absence (Section 9)
"Solely": human involvement must be meaningful	The reviewer can and does change outcomes	Draft & Approve; Readiness Condition 6; score-card metrics 8 and 9 show that review is not nominal

OBLIGATION	WHAT IT DEMANDS OF A DELEGATION DESIGN	VAOM MECHANISM AND EVIDENCE
Article 22(3): human intervention, expressing a view, contesting	A path to a human with authority	Layer 7 escalation and override paths; the accountable role in the matrix; override logs in Layer 6
Article 15(1)(h): meaningful information about the logic involved	Explain the procedure and principles actually applied to this decision	Decision traces (Section 9); the per-decision dimension breakdown and policy version (Section 8); traceability identifiers (Section 12)
Article 22(4): special category data	A stricter basis	Dimension 3 routing; Layer 1 detection of personal and special category data

The mapping holds under the current text and under the proposal: both require that the organization know which decisions are solely automated, justify them, and provide human intervention and explanation on request.

DORA, NIS2, and financial supervision

DORA's contractual requirements for ICT services (Article 30) are the closest the EU comes to requiring notice of vendor model changes. For services supporting critical or important functions, Article 30(3)(b) requires contracts to include the provider's obligation to notify "any development that might have a material impact" on its ability to deliver.¹⁷ There is no general duty to announce model version changes. Section 12's contractual controls should be read against that narrower baseline.

The European Supervisory Authorities' joint statement JC 2026 25 (31 July 2026)¹⁸ addresses ICT risk from frontier AI models, mainly AI-enabled attacks, rather than the governance of firms' own agents. Its relevance here is its insistence that "management bodies must ensure that accountability keeps pace with emerging risks" and that DORA's framework remains the vehicle. EIOPA's Opinion on AI governance and risk management (EIOPA-BoS-25-360, 6 August 2025),¹⁹ and BaFin's non-binding guidance on ICT risks in the use of AI (December 2025),²⁰ treat AI within existing governance and DORA structures rather than as a special case.

NIS2 is subject to a Commission proposal for targeted amendments (COM(2026) 13, 20 January 2026)²¹, published alongside a proposed new Cybersecurity Act. Both are proposals.

A 2

United Kingdom

The UK has diverged from the EU on automated decision-making. Section 80 of the Data (Use and Access) Act 2025, fully in force since 5 February 2026,²² replaced Article 22 of the UK GDPR with Articles 22A to 22D. The new regime permits solely automated decisions subject to safeguards: information about the decision, the ability to make representations, human intervention, and the ability to contest. Restrictions remain for special category data. A decision is "based solely on automated processing if there is no meaningful human involvement".

The ICO's draft guidance on the new regime (consultation 31 March to 29 May 2026, not yet final) sets out what meaningful involvement requires: a human who reviews "at an appropriate point to ensure actual impact on the outcome", has "discretion and authority to alter the decision", is "suitably trained and qualified", and takes account of the relevant data. It adds that "ad hoc spot-checking isn't sufficient".²³ These are close to a specification of VAOM's Band B reviewer and of Readiness Condition 6.

In financial services, the FCA's Mills Review of AI in retail financial services (6 July 2026) recommends that the FCA "enable the foundations for agentic finance".²⁴ In responding, the FCA's Chair described its approach to AI as "relying on the Consumer Duty and Senior Managers Regime". The Senior Managers and Certification Regime is the UK precedent for Readiness Condition 5: each prescribed responsibility is normally held by one person, and sharing is permitted only where it is "appropriate and justifiable", in which case the holders are jointly accountable.²⁵ The PRA's supervisory statement SS28/15 (April 2026 version) treats the regime as "consistent with the principle of collective decision-making" while holding senior managers accountable for "their individual contributions to collective decisions",²⁶ which is the distinction VAOM draws between consent rights and accountability.

A 3

United States

Model risk management

On 17 April 2026 the Federal Reserve (SR 26–2)²⁷ and the OCC (Bulletin 2026–13)²⁸ issued revised guidance on model risk management, superseding SR 11–7 and its OCC counterpart. The guidance excludes generative and agentic AI explicitly: "Generative AI and agentic AI models are novel and rapidly evolving. As such, they are not within the scope of this guidance. Nonetheless, a banking organization's risk management and governance practices should guide the determination of appropriate governance and controls for any tools, processes, or systems not covered in this document." For US banks, the governance of agentic systems is therefore left to the institution's own practices, which is the space a method such as VAOM is designed to fill. References to SR 11–7 as current guidance are now out of date.

NIST

NIST's AI Agent Standards Initiative, launched on 17 February 2026 and coordinated by the Center for AI Standards and Innovation, had published no standards deliverables as of September 2026.²⁹ Its activity so far comprises a request for information on agent security, listening sessions, and a National Cybersecurity Center of Excellence concept paper, *Accelerating the Adoption of Software and Artificial Intelligence Agent Identity and Authorization* (initial public draft, 5 February 2026),³⁰ which is the likely precursor of a NIST practice guide on agent identity.

The SP 800–53 control overlays for securing AI systems (COSaIS) remain pre-final. Only the predictive AI overlay has a document, an annotated outline of 8 January 2026; the single-agent and multi-agent overlays are at concept stage.³¹ VAOM's commitment stands: a control-by-control mapping of VAOM artifacts to the agent overlays will be published in this annex when the first public draft appears. NIST IR 8596, the Cyber AI Profile of the Cybersecurity Framework, is at initial preliminary draft (16 December 2025).³²

State law

Two state regimes create rights that map to VAOM's human review and contest paths:

REGIME	EFFECTIVE	WHAT IT REQUIRES	VAOM MAPPING
Colorado SB 26–189, <i>Automated Decision-Making Technology</i> (re-peals and re-enacts the Colorado AI Act of SB 24–205) ³³	Consequential decisions on or after 1 January 2027	After an adverse decision materially influenced by the technology, on request, "an opportunity for meaningful human review and reconsideration ... to the extent commercially reasonable". Meaningful human review is review by someone "who has authority to approve, modify, or override" the decision, who considers the available evidence and is trained to review	Layer 7 contest path; the Band B reviewer definition; Readiness Condition 6
California CPPA regulations on automated decisionmaking technology ³⁴	ADMT compliance from 1 January 2027	An opt-out from ADMT for significant decisions, unless the business instead offers an appeal to a human reviewer "who has the authority to overturn the decision"	The same, with the appeal path as an alternative to opt-out

Texas's Responsible Artificial Intelligence Governance Act (in force 1 January 2026)³⁵ is largely a regime of prohibited uses enforced by the Attorney General and does not create a comparable human review right.

A 4

Asia-Pacific and International

Singapore

IMDA Model AI Governance Framework for Agentic AI (version 1.0, January 2026; version 1.5, 20 May 2026, updated 5 June 2026). The first government framework specifically for AI agents; its structure parallels VAOM's architecture dimension by dimension. Version 1.5 added multi-agent systemic-risk factors and practices against automation bias, recommending that organizations track "human override rate", where "a low rate may signal rubber-stamping behaviours", and "human response times during review", where "a shorter time may signal automation bias or review fatigue".³⁶ These are VAOM's review engagement metric (Section 11, metric 8).

IMDA DIMENSION	VAOM MECHANISM
Assess and bound risks upfront	Delegation Discovery & Design (Section 5); Authority Decomposition
Make humans meaningfully accountable	Readiness Conditions 5 and 6; accountable roles; Layer 7
Technical controls and processes	Confidence Gate; identity-bound scopes (Section 9); Continuous Assurance (Section 10)
Tiered autonomy (fully automated / human-approved / off-limits)	Bands A / B / non-delegable, with the added confidence dimension and decision-level granularity

IMDA, Legal Responsibility for AI Agents (discussion paper, May 2026). A working group of Singapore's legal community concluded that many cases "may be capable of being addressed through the common law, such as contract and the tort of negligence, though the law may need to be adapted", and identified who bears "unforeseeable losses" from agent actions as the hardest open question.³⁷ See A8.

MAS, Safeguards for Agentic Finance at Runtime (SAFR) (version 1.0, July 2026). An industry white paper developed under MAS's BuildFin.ai initiative, not regulatory guidance. It is the closest external counterpart to VAOM's runtime structure published to date. Four runtime components (Agent Identity, a Controls Repository, a Disposition Engine, and an append-only Audit Log) resolve every proposed agent action to one of four outcomes:³⁸

SAFR DISPOSITION	DEFINITION (ABRIDGED FROM SAFR)	VAOM COUNTERPART
Deny	Violates a hard regulatory or policy constraint, or presents a risk profile above defined thresholds; rejected before execution	Non-delegable; hard floors; tool or scope absence
Escalate	Within scope and below hard constraints, but above the threshold for autonomous execution; held for human review	Band B (and Band C where the reviewer is a specialist)
Auto-Execute	Within scope, below hard constraints, within defined risk thresholds	Band A
Observe	Permitted to proceed but flagged for monitoring and later review	Execute & Audit sampling; shadow and pilot periods; behavioral envelope alerts

SAFR's Controls Repository draws on "organisational policies, regulatory requirements, product rules, and user-provided mandates", and it notes that "the calibration is set at design time". Both are the Delegation Authority Matrix under another name. The relationship is the one VAOM has with every runtime engine: SAFR specifies how a disposition is reached and recorded; VAOM is the method that produces the controls and calibration it evaluates against. MAS's own Guidelines on AI Risk Management, consulted on from November 2025 and stated to cover agentic AI, were not final as of September 2026.³⁹

China

The *Implementation Opinions on the Standardized Application and Innovative Development of Intelligent Agents*, issued jointly by the Cyberspace Administration of China, the National Development and Reform Commission, and the Ministry of Industry and Information Technology on 8 May 2026, direct that the boundaries and required permissions be clarified⁴⁰ for three decision modes: decisions reserved to the user personally, decisions requiring the user's authorization, and autonomous decisions by the agent. Users are to keep "the right to know and final decision-making power" over autonomous decisions, and agent operations "must not exceed the scope of the user's authorization". The three modes correspond to non-delegable, Band B, and Band A within scope. The document is policy guidance, not a statute with penalties, and it pairs the modes with graded governance by scenario, including filing and testing in sensitive sectors.

Korea

Korea's AI Basic Act has been in force since 22 January 2026.⁴¹ Operators of high-impact AI must implement risk management, explanation to the extent technically feasible, user protection, human management and supervision, and documentation (Article 34). Administrative fines of up to KRW 30 million attach to failures of transparency notice, of designating a domestic representative, and of complying with suspension or correction orders, not to the Article 34 duties themselves. The Ministry of Science and ICT has announced a guidance period of at least one year before fines are imposed.

OECD, Council of Europe, and ISO/IEC

The OECD's *The Agentic AI Landscape and Its Conceptual Foundations* (AI Paper No. 56, February 2026)⁴² sets out the policy vocabulary. Its *Agentic AI in Organisations: Early Insights from Practitioner Interviews* (AI Paper No. 65, September 2026),⁴³ based on interviews with 25 organisations, found checkpoint-based oversight emerging: "autonomous execution within defined boundaries, while requiring human review for high-impact, potentially irreversible actions", with continuous human-in-the-loop review "widely seen as impractical at scale". That is the Delegation Authority Matrix described from the practitioner side.

The Council of Europe's Framework Convention on Artificial Intelligence (CETS No. 225) was ratified by the European Union in May 2026 and was not yet in force as of September 2026.⁴⁴

Among ISO/IEC standards, ISO/IEC 42001 (AI management systems)⁴⁵ remains the management-system frame that VAOM's Layer 6 evidence supports. ISO/IEC 42005:2025 (AI system impact assessment),⁴⁶ complements the Delegation Readiness Assessment, and ISO/IEC 42006:2025 sets requirements for bodies certifying AI management systems.⁴⁷ A second amendment to ISO/IEC 22989 (AI concepts and terminology), expected to address agent terminology, reached committee-draft stage in August 2026.⁴⁸

A 5

Runtime Control and Agent Security Standards

OWASP threat taxonomies. The *OWASP Top 10 for Agentic Applications for 2026* (December 2025)⁴⁹ catalogs what can go wrong; VAOM's structures are among the controls. Agent identity and privilege abuse is countered by delegated execution contexts and scope compilation (Section 9); excessive autonomy by the matrix and band ceilings; multi-agent exploitation by attenuation, chain limits, and the delegation-laundering defenses of Section 5.3. The companion *OWASP Top 10 for LLM Applications 2026* was published on 3 August 2026.⁵⁰ Security teams should use both as an adversarial test suite for a VAOM implementation: every entry is a question the delegation design should already answer, and boundary verification (Section 9) is where the answers are tested.

OWASP Agent Control Standard (ACS). A vendor-neutral specification for the runtime control plane, adopted into the OWASP GenAI Security Project in September 2026.⁵¹ The specification is at version 0.1.0 (repository releases 0.1.1 on 11 August 2026 and 0.1.2 on 21 September 2026), and its reference implementation is described as a proof of concept. It instruments agents at defined lifecycle hooks (user messages, tool call requests and results, knowledge and memory retrieval, memory writes, sub-agent start and stop, skill loading, and session and turn boundaries), returns inline allow, deny, modify, ask, or defer verdicts through a Guardian Agent pattern whose deterministic policy layer always runs first, and standardizes telemetry and

agent bills of materials on open conventions. An Identity for Agents workstream was added in September 2026. ACS standardizes how a boundary is enforced and observed; VAOM produces the boundary. Matrix rows and circuit-breaker conditions compile to intervention-point policies, and a VAOM guardian (Section 10) has the ACS Guardian Agent as a natural implementation target, with VAOM supplying what the standard presupposes: the guardian's own matrix row, its accountable role, its containment-only authority, and its lease. Microsoft's Agent Governance Toolkit separately uses "Agent Control Specification", also abbreviated ACS, for its own policy decision runtime; ⁵² references here are to the OWASP standard.

Temporal policy languages. AWS's Dogwood (August 2026), ⁵³ an open extension of Cedar integrated with Amazon Bedrock AgentCore Policy and grounded in metric first-order temporal logic, was the first runtime enforcement language for agent action sequences from a major provider. It is the compilation target for the temporal path of Section 9:

VAOM MECHANISM	TEMPORAL POLICY IMPLEMENTATION
Band B human review	Prerequisite approval event within a window
Cumulative value ceilings (Section 5.5)	Windowed sum caps
Frequency and volume limits; circuit breakers	Rate rules
Delegation Authority Matrix row	The policy someone must first design

AgentCore's accompanying capabilities include session-level cumulative spend enforcement and permissions that "can narrow automatically when a person is no longer engaged", ⁵⁴ a direct compile target for the transition from Band B to autonomous operation. Two caveats carry over from the provider's guidance: temporal enforcement depends on a complete, authenticated event history, and the open reference interpreter is for exploration, with production enforcement in the hosted service. Temporal operators compile the deterministic slice of assurance (counts, sums, sequence prerequisites). They do not subsume the calibration and anomaly work of Sections 8, 10, and 11: a rate rule can refuse the eleventh call in a window, but it cannot notice that a verifier has stopped disagreeing.

Gateways and guardians. Agent Router, formerly Envoy AI Gateway, joined the Agentic AI Foundation in September 2026 ⁵⁵ as an open-source routing and enforcement point for tool calls, with credentials and quotas declared once. Commercial guardian products reached the market in the same period, among them Cyera's Agent Guardian (3 August 2026) ⁵⁶ and CrowdStrike's Falcon Guardian (1 September 2026) ⁵⁷. Gartner's *Market Guide for Guardian Agents* (February 2026) ⁵⁸ frames the category. VAOM's clarification applies to every product in it: a guardian is itself a delegated agent, with a matrix row, an accountable role, containment-only authority, and human oversight of its own false-suspension and missed-detection rates.

Runtime evidence. TRACE (Trust, Runtime Attestation and Compliance Evidence), contributed to the Linux Foundation by OPAQUE and developed with AMD, Intel, Microsoft, and TII (August 2026), ⁵⁹ produces hardware-enforced governance records binding runtime, software, policies, and tool use. It addresses the caveat of Section 9 that temporal enforcement is only as trustworthy as its event history, and is a candidate substrate for the decision records of Section 12.

Identity, Authorization, and Agent Protocols

IETF. The WIMSE working group adopted *AI Identity Management System* (draft-ietf-wimse-aims-00, 15 September 2026),⁶⁰ which composes existing mechanisms (workload identifiers, short-lived credentials, transaction tokens, token exchange, identity chaining, shared signals for dynamic authorization changes, and client-initiated backchannel authentication for human involvement) into an identity architecture for agents. Its treatment of policy states VAOM's position from the standards side. Because authorization parameters are "highly deployment and risk-model-specific (and often reflect local governance, regulatory, and operational constraints), the policy model and document format are out of scope for this framework and are not recommended as a target for standardization"; implementations may use any policy format "provided it is versioned, reviewable, and supports consistent evaluation". The Delegation Authority Matrix is such a policy, and Section 5 is the method for writing it. The draft also notes that confirming actions with the user mid-execution "may" need further specification work.

Related work: *Transaction Tokens* (draft-ietf-oauth-transaction-tokens-11, July 2026, at working group consensus)⁶¹ is the closest standard carrier for the task-bound delegated execution context within a trust domain. The individual drafts on attenuating authorization tokens for agentic delegation chains,⁶² and an OAuth actor profile for delegation,⁶³ give the attenuation rule and the accountable-role-plus-agent audit pair native token semantics. The OpenID AuthZEN Authorization API 1.0, a Final Specification since January 2026,⁶⁴ standardizes the decision interface through which scopes are evaluated. OpenID's Continuous Access Evaluation Profile and Shared Signals Framework (Final, September 2025),⁶⁵ carry the revocation signals used by the delegation lease. An agent communication protocols working group (agentproto) has been proposed at the IETF since 9 September 2026.⁶⁶

Model Context Protocol. The 2026-07-28 revision of MCP moved to a stateless core⁶⁷ and hardened authorization (issuer validation, a move away from dynamic client registration toward client metadata documents).⁶⁸ Three changes map to VAOM:

- *Enterprise-Managed Authorization*, a stable extension,⁶⁹ lets the corporate identity provider decide which tool servers an agent may reach. It is a compile target for tool absence enforced at the identity layer rather than in the agent's own configuration.
- Multi round-trip requests let a tool request human input in the middle of a call, and the Tasks extension surfaces an `input_required` state for long-running work.⁷⁰ Both give the Band B approval event a protocol-level carrier.
- When a client steps up scopes, the specification has it request the union of existing and new scopes. Step-up therefore only widens authority, and VAOM treats every step-up as a decision the matrix must authorize rather than a protocol convenience.

Agent-to-agent and the Agentic AI Foundation. The A2A protocol (version 1.0, March 2026, with signed agent cards) joined the Agentic AI Foundation on 27 August 2026,⁷¹ alongside MCP and Agent Router. Signed agent cards are a verifiable anchor for the registered, versioned agent identity of Section 9 in chains that cross organizational boundaries.

Identity platforms. Microsoft Entra Agent ID reached general availability in April 2026,⁷² and Microsoft Agent 365 on 1 May 2026.⁷⁴ Entra requires at least one *sponsor* for every agent identity: "business representatives accountable for the agent's purpose and lifecycle decisions", who may authorize suspension during incid-

ents, with sponsorship transferred automatically to the sponsor's manager when the sponsor leaves, so that "there's always a human user accountable".⁷³ The sponsor is a product implementation of the accountable role of Readiness Condition 5. The transfer rule is a platform-level answer to the orphaned-delegation problem; the matrix entry should still name the re-attestation deputies (Section 9) explicitly, one for each re-attestor.

A 7

Agent Payments

Worked example 7E describes delegated purchasing in protocol-neutral terms. The protocols behind it, as of September 2026:

PROTOCOL	STATUS	HOW IT EXPRESSES BOUNDED DELEGATED AUTHORITY	VAOM MAPPING
Google, Agent Payments Protocol (AP2) ⁷⁵	v0.2.0, 28 April 2026; donated to the FIDO Alliance	Open mandates carry constraints; closed mandates fix final values. With the human present, the user approves closed mandates. With the human not present, the user approves open mandates and the agent signs closed mandates bound to them	Human-present closed mandate: Band B approval record. Open mandate: Band A within constraints, a lease-like grant. Agent-signed closed mandate: evidence of compliance, not of review
Mastercard, Verifiable Intent ⁷⁶	Draft v0.1 (February 2026); contributed to the FIDO Alliance	Layered credentials: an issuer-bound user credential, a user mandate with constraints (amount range, line items, approved merchants), and, in autonomous mode only, short-lived agent credentials checked against the mandate	The attenuation rule enforced by the counterparty; the short-lived layer is a response context
OpenAI and Stripe, Agentic Commerce Protocol, Delegate Payment API ⁷⁷	Version 2026-04-17, beta	A single-use allowance: maximum amount, currency, merchant, checkout session, expiry	A task-bound delegated execution context for one purchase
Visa, Trusted Agent Protocol ⁷⁸	Launched 14 October 2025; public repository last updated October 2025	A framework for agent-to-merchant trust signals rather than explicit spending limits	Agent identity verification at the counterparty

The FIDO Alliance formed an Agentic Authentication Technical Working Group and a Payments Technical Working Group on 28 April 2026, whose remit includes agents acting within defined limits.⁷⁹ No draft specification had been published as of September 2026.

A 8

Liability and Insurance

Product liability. The revised Product Liability Directive, (EU) 2024/2853, applies to products placed on the market or put into service after 9 December 2026.⁸⁰ Software is a product. Defectiveness takes account of "the effect on the product of any ability to continue to learn or acquire new features after it is placed on the market", and the manufacturer is not exempt where a defect is due to software updates or upgrades within its control, to a missing update needed to maintain safety, or to a substantial modification, which can include one brought about by an AI system's continuous learning. For organizations that build or substantially modify agentic products, Section 12's channels of change are therefore also channels of liability, and its registries and traceability are the evidence of what changed and when.

Fault liability. The Commission withdrew its proposed AI Liability Directive (COM(2022) 496); the withdrawal is recorded on 6 October 2025.⁸¹ There is no EU-level AI-specific fault liability regime, so fault liability for agent actions remains a matter of national law. IMDA's discussion paper reaches a similar position for Singapore: negligence and contract can address many cases, with difficulty locating fault, the standard of care, and remoteness, and with "unforeseeable losses" the hardest case. In a negligence analysis, VAOM's artifacts (the matrix with its derivations, the readiness evidence, boundary verification results, the scorecard) are the record of the care the organization took.

Insurance. The market is dividing between exclusion and affirmative cover. Trade press reports generative AI exclusion endorsements for commercial general liability from ISO (Verisk), including CG 40 47, CG 40 48, and CG 35 08, and an absolute AI exclusion introduced by Berkley on management and professional liability lines; we have not seen the filings themselves.⁸² Affirmative AI liability products are offered, among others, by Armilla (underwritten at Lloyd's, with Chaucer as a partner),⁸³ Munich Re (aiSure),⁸⁴ HSB (AI Liability Insurance),⁸⁵ and Testudo, a Lloyd's coverholder.⁸⁶ Certification is emerging as the bridge between the two: the AIUC-1 standard for AI agents is revised quarterly (latest release 15 July 2026,⁸⁷ which added requirements for code-generating agents on secrets leakage and secure code, and restructured its agent identity and access controls), and certifications against it have been announced, among them ElevenLabs in February 2026.⁸⁸ A crosswalk from VAOM artifacts to AIUC-1 requirements is planned for a future revision of this annex, once it can be built control by control against the current release.

A 9

Research Alignment

Academic work on agent autonomy, notably the five-level framework of Feng, McDonald, and Zhang,⁸⁹ characterizes autonomy by the role of the user (operator, collaborator, consultant, approver, observer) and arrives at a structure recognizably parallel to VAOM's delegation patterns: Prepare & Present places the human as operator, Draft & Approve as approver, and Execute & Audit and Monitor & Intervene as observer. Its proposal of "autonomy certificates", formal statements of the maximum autonomy an agent may operate at, has a natural substrate in a signed Delegation Authority Matrix entry compiled into scopes: an autonomy certifi-

ate with enforcement attached. The older levels-of-automation literature, in particular Parasuraman, Sheridan, and Wickens (2000) and Bainbridge's *Ironies of Automation* (1983), is acknowledged in the whitepaper itself (Section 5.3).

Annex References

Listed by section. URLs were checked on 27 September 2026. Where only a secondary source was available, the entry says so.

1. European Parliament and Council of the European Union. *Regulation (EU) 2026/1744 of 8 July 2026 amending Regulations (EU) 2024/1689, (EU) 2018/1139 and (EU) 2023/1230 as regards the simplification of the implementation of harmonised rules on artificial intelligence (Digital Omnibus on AI)*. Regulation; OJ L 24 July 2026; in force 27 July 2026. <https://eur-lex.europa.eu/eli/reg/2026/1744/oj> [Regulatory text]
2. European Parliament and Council of the European Union. *Regulation (EU) 2024/1689 laying down harmonised rules on artificial intelligence (Artificial Intelligence Act)*. Regulation, 13 June 2024; as amended by Regulation (EU) 2026/1744. <https://eur-lex.europa.eu/eli/reg/2024/1689/oj> [Regulatory text]
3. European Commission, AI Act Service Desk. *How are AI agents addressed within the AI Act?* FAQ, undated (accessed September 2026). <https://ai-act-service-desk.ec.europa.eu/en/ai-act/faq/how-are-ai-agents-addressed-within-ai-act-0> [Framework guidance]
4. European Commission. *Draft Commission guidelines on the classification of high-risk AI systems*. Draft for targeted consultation, 19 May 2026. <https://digital-strategy.ec.europa.eu/en/library/draft-commission-guidelines-classification-high-risk-ai-systems> [Framework guidance (draft)]
5. European Commission. *Guidelines on transparency obligations for providers and deployers of AI systems*. Commission guidelines, 20 July 2026. <https://digital-strategy.ec.europa.eu/en/library/guidelines-transparency-obligations-providers-and-deployers-ai-systems> [Framework guidance]
6. European Commission. *Commission publishes Code of Practice on marking and labelling AI-generated content*. News release, 10 June 2026. <https://digital-strategy.ec.europa.eu/en/news/commission-publishes-code-practice-marking-and-labelling-ai-generated-content> [Framework guidance]
7. CEN-CENELEC. *EN 18286 in the Spotlight: Supporting Compliance with the AI Act*. News release, 31 July 2026. <https://www.cencenelec.eu/news-events/news/2026/en-in-the-spotlight/2026-07-30-ai-quality-management/> [Standard]
8. CEN-CENELEC JTC 21. *prEN 18228 Artificial intelligence: Risk management*. Draft European standard, CEN Enquiry 2026 (status from secondary sources). <https://standardsdevelopment.bsigroup.com/projects/2025-01990> [Standard (draft)]
9. CEN-CENELEC JTC 21. *prEN 18229-3 AI trustworthiness framework, Part 3: Human oversight*. Draft European standard, CEN Enquiry, comment period 3 August to 7 October 2026. <https://kommentera.sis.se/Home/Details/15712> [Standard (draft)]
10. Adam Leon Smith. *prEN 18229-3 reaches Enquiry*. Substack post, 30 July 2026 (commentary on a non-public draft). <https://adamleonsmith.substack.com/p/pren-18229-3-reaches-enquiry> [Framework guidance (commentary)]
11. Court of Justice of the European Union. *Judgment of 7 December 2023, OQ v Land Hessen (SCHUFA Holding)*, Case C-634/21. ECLI:EU:C:2023:957. <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:62021CJ0634> [Regulatory text (case law)]
12. Court of Justice of the European Union. *Judgment of 27 February 2025, CK v Magistrat der Stadt Wien (Dun & Bradstreet Austria)*, Case C-203/22. <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:62022CJ0203> [Regulatory text (case law)]
13. Article 29 Data Protection Working Party. *Guidelines on Automated individual decision-making and Profiling for the purposes of Regulation 2016/679 (WP251rev.01)*. Adopted 3 October 2017, revised 6 February 2018; endorsed by the EDPB 25 May 2018. <https://ec.europa.eu/newsroom/article29/items/612053> [Framework guidance]
14. European Commission. *Proposal for a Regulation ... as regards the simplification of the digital legislative framework (Digital Omnibus)*. COM(2025) 837 final, 19 November 2025 (proposal; procedure 2025/0360(COD)). <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:52025PC0837> [Regulatory text (proposal)]

15. European Parliament. *Legislative Observatory: 2025/0360(COD) Digital Omnibus*. Procedure file, accessed 27 September 2026. [https://oeil.secure.europarl.europa.eu/oeil/en/procedure-file?reference=2025/0360\(COD\)](https://oeil.secure.europarl.europa.eu/oeil/en/procedure-file?reference=2025/0360(COD)) [Regulatory text (procedure)]
16. European Data Protection Board and European Data Protection Supervisor. *EDPB-EDPS Joint Opinion 2/2026 on the Proposal for a Regulation as regards the simplification of the digital legislative framework (Digital Omnibus)*. Adopted 10 February 2026. https://www.edpb.europa.eu/system/files/documents/2026-02/edpb_edps_jointopinion_202602_digitalomnibus_en.pdf [Framework guidance]
17. European Parliament and Council of the European Union. *Regulation (EU) 2022/2554 on digital operational resilience for the financial sector (DORA)*. Regulation, 14 December 2022; applies from 17 January 2025. <https://eur-lex.europa.eu/eli/reg/2022/2554/oj> [Regulatory text]
18. European Supervisory Authorities (EBA, EIOPA, ESMA). *ESA Statement: Toward a consistent and risk-based approach for ICT risks from frontier AI models (JC 2026 25)*. Joint statement, 31 July 2026. https://www.esma.europa.eu/sites/default/files/2026-07/JC_2026_25_ESA_statement_on_frontier_AI_models.pdf [Framework guidance]
19. European Insurance and Occupational Pensions Authority. *Opinion on Artificial Intelligence Governance and Risk Management (EIOPA-BoS-25-360)*. Supervisory opinion, 6 August 2025. https://www.eiopa.europa.eu/publications/opinion-artificial-intelligence-governance-and-risk-management_en [Framework guidance]
20. Bundesanstalt für Finanzdienstleistungsaufsicht (BaFin). *Guidance on ICT Risks in the Use of Artificial Intelligence at Financial Entities*. Non-binding guidance, December 2025. https://www.bafin.de/SharedDocs/Downloads/EN/Anlage/dl_Anlage_orientierungshilfe_IKT_Risiken_bei_KI_en.html [Framework guidance]
21. European Commission. *Proposal for a Directive as regards simplification measures and alignment with the Cybersecurity Act (targeted amendments to NIS2)*. COM(2026) 13, 20 January 2026 (proposal). <https://digital-strategy.ec.europa.eu/en/library/proposal-directive-regards-simplification-measures-and-alignment-cybersecurity-act> [Regulatory text (proposal)]
22. UK Parliament. *Data (Use and Access) Act 2025, section 80, inserting UK GDPR Articles 22A to 22D*. Act of Parliament, 2025 c. 18; s.80 fully in force 5 February 2026. <https://www.legislation.gov.uk/ukpga/2025/18/section/80> [Regulatory text]
23. Information Commissioner's Office. *Automated decision-making, including profiling*. Draft updated guidance, 31 March 2026 (consultation closed 29 May 2026; not final). <https://ico.org.uk/for-organisations/uk-gdpr-guidance-and-resources/individual-rights/automated-decision-making/what-does-the-uk-gdpr-say-about-adm/> [Framework guidance (draft)]
24. Financial Conduct Authority (Sheldon Mills). *The Mills Review: AI and the future of retail financial services*. Review and press release, 6 July 2026. <https://www.fca.org.uk/news/press-releases/fca-publishes-landmark-review-impact-ai-retail-financial-services> [Framework guidance]
25. Financial Conduct Authority. *The Senior Managers and Certification Regime: Guide for FCA solo-regulated firms*. Regulatory guide, updated July 2019. <https://www.fca.org.uk/publication/policy/guide-for-fca-solo-regulated-firms.pdf> [Framework guidance]
26. Prudential Regulation Authority. *Supervisory Statement SS28/15 Strengthening individual accountability in banking*. April 2026 update (effective 24 April 2026). <https://www.bankofengland.co.uk/prudential-regulation/publication/2015/strengthening-individual-accountability-in-banking-ss> [Framework guidance]
27. Board of Governors of the Federal Reserve System. *SR 26-2: Revised Guidance on Model Risk Management*. Supervisory letter, 17 April 2026. <https://www.federalreserve.gov/supervisionreg/srletters/SR2602.htm> [Regulatory text]
28. Office of the Comptroller of the Currency. *Model Risk Management: Revised Guidance*. OCC Bulletin 2026-13, 17 April 2026. <https://www.occ.gov/news-issuances/bulletins/2026/bulletin-2026-13.html> [Regulatory text]
29. NIST Center for AI Standards and Innovation. *AI Agent Standards Initiative*. Launched 17 February 2026. <https://www.nist.gov/artificial-intelligence/ai-agent-standards-initiative> [Standard (programme)]
30. NIST National Cybersecurity Center of Excellence (Harold Booth, William Fisher, Ryan Galluzzo, Joshua Roberts). *Accelerating the Adoption of Software and Artificial Intelligence Agent Identity and Authorization*. Concept paper, initial public draft, 5 February 2026. <https://csrc.nist.gov/pubs/other/2026/02/05/accelerating-the-adoption-of-software-and-ai-agent/ipd> [Standard (draft)]
31. NIST. *SP 800-53 Control Overlays for Securing AI Systems (COSAiS)*. Project page; predictive AI overlay annotated outline, 8 January 2026; agent overlays at concept stage. <https://csrc.nist.gov/projects/cosais> [Standard (pre-final)]
32. NIST. *NIST IR 8596: Cybersecurity Framework Profile for Artificial Intelligence (Cyber AI Profile)*. Initial preliminary draft, 16 December 2025. <https://csrc.nist.gov/pubs/ir/8596/iprd> [Standard (draft)]
33. Colorado General Assembly. *SB 26-189, Automated Decision-Making Technology*. Signed 14 May 2026; effective 1 January 2027. <https://leg.colorado.gov/bills/sb26-189> [Regulatory text]

34. California Privacy Protection Agency. *CCPA Updates, Cybersecurity Audits, Risk Assessments, Automated Decisionmaking Technology (ADMT), and Insurance Regulations*. Approved 23 September 2025; ADMT compliance from 1 January 2027. https://coppa.ca.gov/regulations/ccpa_updates.html [Regulatory text]
35. Texas Legislature. *H.B. 149, Texas Responsible Artificial Intelligence Governance Act*. Enrolled act; effective 1 January 2026. <https://capitol.texas.gov/tlodocs/89R/billtext/pdf/HB00149F.pdf> [Regulatory text]
36. Infocomm Media Development Authority, Singapore. *Model AI Governance Framework for Agentic AI*. Version 1.0, 22 January 2026; version 1.5, 20 May 2026 (updated 5 June 2026) (§2.2.2: human override rate and review response time). <https://www.imda.gov.sg/-/media/imda/files/about/emerging-tech-and-research/artificial-intelligence/mgf-for-agentic-ai.pdf> [Framework guidance]
37. Infocomm Media Development Authority, Singapore. *Legal Responsibility for AI Agents*. Discussion paper, May 2026. <https://www.imda.gov.sg/-/media/imda/files/about/emerging-tech-and-research/artificial-intelligence/agents-legal-responsibility.pdf> [Framework guidance]
38. Monetary Authority of Singapore and industry contributors (BuildFin.ai). *Safeguards for Agentic Finance at Runtime (SAFR)*. Industry white paper, version 1.0, July 2026 (not regulatory guidance). <https://www.mas.gov.sg/-/media/mas-media-library/development/fintech/ai-safr/safr.pdf> [Framework guidance]
39. Monetary Authority of Singapore. *Consultation Paper on Proposed Guidelines on Artificial Intelligence Risk Management*. 13 November 2025 (not finalised as of September 2026). <https://www.mas.gov.sg/publications/consultations/2025/consultation-paper-on-guidelines-on-artificial-intelligence-risk-management> [Regulatory text (consultation)]
40. Cyberspace Administration of China, National Development and Reform Commission, and Ministry of Industry and Information Technology. *Implementation Opinions on the Standardized Application and Innovative Development of Intelligent Agents* (智能体规范应用与创新发展实施意见). Policy guidance, 8 May 2026 (item 6). https://www.cac.gov.cn/2026-05/08/c_1779979789523320.htm [Regulatory text (policy guidance)]
41. Republic of Korea. *Framework Act on the Development of Artificial Intelligence and the Establishment of a Foundation for Trust (AI Basic Act)*. Act No. 20676, as amended; in force 22 January 2026. <https://www.law.go.kr/법령/인공지능발전과신뢰기반조성등에관한기본법> [Regulatory text]
42. OECD (Luis Aranda and Kasumi Sugimoto). *The Agentic AI Landscape and Its Conceptual Foundations*. OECD Artificial Intelligence Papers No. 56, February 2026. <https://doi.org/10.1787/396cf758-en> [Framework guidance]
43. OECD (Merve Hickok, Rentaro Iida, Luis Aranda, and Kasumi Sugimoto). *Agentic AI in Organisations: Early Insights from Practitioner Interviews*. OECD Artificial Intelligence Papers No. 65, 16 September 2026. <https://doi.org/10.1787/1257a26f-en> [Survey]
44. Council of Europe. *Framework Convention on Artificial Intelligence and Human Rights, Democracy and the Rule of Law* (CETS No. 225). Opened for signature 5 September 2024; not in force as of September 2026 (status from secondary sources). <https://www.coe.int/en/web/conventions/full-list?module=treaty-detail&treaty-num=225> [Regulatory text]
45. ISO/IEC. *ISO/IEC 42001:2023 Information technology: Artificial intelligence: Management system*. International Standard, 2023. <https://www.iso.org/standard/42001> [Standard]
46. ISO/IEC. *ISO/IEC 42005:2025 Information technology: Artificial intelligence: AI system impact assessment*. International Standard, 2025. <https://www.iso.org/standard/42005> [Standard]
47. ISO/IEC. *ISO/IEC 42006:2025 Information technology: Artificial intelligence: Requirements for bodies providing audit and certification of artificial intelligence management systems*. International Standard, July 2025. <https://www.iso.org/standard/42006> [Standard]
48. ISO/IEC JTC 1/SC 42. *ISO/IEC 22989:2022/CD Amd 2 Artificial intelligence concepts and terminology, Amendment 2*. Committee draft, registered August 2026. <https://www.iso.org/standard/93144.html> [Standard (draft)]
49. OWASP GenAI Security Project. *OWASP Top 10 for Agentic Applications for 2026*. Published 9 December 2025. <https://genai.owasp.org/resource/owasp-top-10-for-agentic-applications-for-2026/> [Framework guidance]
50. OWASP GenAI Security Project. *OWASP Top 10 for LLM Applications 2026*. Published 3 August 2026. <https://genai.owasp.org/resource/owasp-genai-llm-top-10-2026/> [Framework guidance]
51. OWASP GenAI Security Project. *Agent Control Standard (ACS)*. Specification v0.1.0 (repository releases v0.1.1, 11 August 2026, and v0.1.2, 21 September 2026); resource page 1 September 2026. <https://genai.owasp.org/resource/agent-control-standard-acs/> [Standard (draft)]
52. Microsoft. *Agent Control Specification (ACS), Agent Governance Toolkit*. Project documentation, 2026 (cited for disambiguation only). <https://microsoft.github.io/agent-governance-toolkit/packages/agent-control-specification/> [Standard (draft)]

53. Amazon Web Services (Marc Brooker, Joseph Tassarotti, and Jean-Baptiste Tristan). *Introducing Dogwood: Runtime Verification for AI Agents*. AWS Open Source Blog, 6 August 2026. <https://aws.amazon.com/blogs/opensource/introducing-dogwood-runtime-verification-for-ai-agents/> [Product release]
54. Amazon Web Services (Madhu Parthasarathy). *Control agent behaviors and cost beyond a single action: new capabilities in Amazon Bedrock AgentCore*. AWS Machine Learning Blog, 6 August 2026. <https://aws.amazon.com/blogs/machine-learning/control-agent-behaviors-and-cost-beyond-a-single-action-new-capabilities-in-amazon-bedrock-agentcore/> [Product release]
55. Agentic AI Foundation. *Envoy AI Gateway becomes Agent Router and joins AAIF*. AAIF blog, 9 September 2026. <https://aaif.io/blog/agent-router-joins-aaif> [Product release]
56. Cyera. *Cyera Launches Agent Guardian to Secure the Autonomous Workforce*. Press release, 3 August 2026. <https://www.cyera.com/press-releases/cyera-launches-agent-guardian-to-secure-the-autonomous-workforce> [Product release]
57. CrowdStrike. *CrowdStrike Unveils Falcon Guardian to Secure AI Agents Where They Execute: On the Endpoint at Runtime*. Press release, 1 September 2026. <https://www.crowdstrike.com/en-us/press-releases/crowdstrike-unveils-falcon-guardian-ai-agent-security/> [Product release]
58. Gartner (Avivah Litan and Daryl Plummer). *Market Guide for Guardian Agents*. Analyst research, February 2026 (paywalled). <https://www.gartner.com/en/documents/7509053> [Forecast]
59. The Linux Foundation. *Linux Foundation Welcomes TRACE to Advance Verifiable Runtime Evidence for AI Workloads*. Press release, 25 August 2026. <https://www.linuxfoundation.org/press/linux-foundation-welcomes-trace-to-advance-verifiable-runtime-evidence-for-ai-workloads> [Standard (programme)]
60. P. Kasselmann, J. Lombardo, Y. Rosomakho, B. Campbell, N. Steele, and A. Parecki. *AI Identity Management System*. IETF Internet-Draft draft-ietf-wimse-aims-00, 15 September 2026 (WIMSE WG document). <https://datatracker.ietf.org/doc/draft-ietf-wimse-aims/> [Standard (draft)]
61. A. Tulshibagwale, G. Fletcher, and P. Kasselmann. *Transaction Tokens*. IETF Internet-Draft draft-ietf-oauth-transaction-tokens-11, 30 July 2026. <https://datatracker.ietf.org/doc/draft-ietf-oauth-transaction-tokens/> [Standard (draft)]
62. N. A. Niyikiza. *Attenuating Authorization Tokens for Agentic Delegation Chains*. IETF Internet-Draft draft-niyikiza-oauth-attenuating-agent-tokens-01, 15 June 2026 (individual submission). <https://datatracker.ietf.org/doc/draft-niyikiza-oauth-attenuating-agent-tokens/> [Standard (draft)]
63. Karl McGuinness. *OAuth Actor Profile for Delegation*. IETF Internet-Draft draft-mcguinness-oauth-actor-profile-00, 30 April 2026 (individual submission). <https://datatracker.ietf.org/doc/draft-mcguinness-oauth-actor-profile/> [Standard (draft)]
64. OpenID Foundation, AuthZEN Working Group. *Authorization API 1.0*. Final Specification, 11 January 2026. https://openid.net/specs/authorization-api-1_0.html [Standard]
65. OpenID Foundation, Shared Signals Working Group. *OpenID Continuous Access Evaluation Profile 1.0 and OpenID Shared Signals Framework Specification 1.0*. Final Specifications, dated 29 August 2025, approved 2 September 2025. https://openid.net/specs/openid-caep-1_0-final.html [Standard]
66. IETF. *Agent Communication Protocols (agentproto)*. Proposed working group since 9 September 2026. <https://datatracker.ietf.org/group/agentproto/about/> [Standard (programme)]
67. Model Context Protocol (David Soria Parra and Den Delimarsky). *The 2026-07-28 Specification*. MCP blog, 28 July 2026. <https://blog.modelcontextprotocol.io/posts/2026-07-28/> [Standard]
68. Model Context Protocol. *Specification, version 2026-07-28: Authorization*. <https://modelcontextprotocol.io/specification/2026-07-28/basic/authorization> [Standard]
69. Model Context Protocol. *Enterprise-Managed Authorization (extension io.modelcontextprotocol/enterprise-managed-authorization)*. Extension specification (stable), 2026. <https://modelcontextprotocol.io/extensions/auth/enterprise-managed-authorization> [Standard]
70. Model Context Protocol. *Tasks (extension io.modelcontextprotocol/tasks)*. Extension documentation, 2026. <https://modelcontextprotocol.io/extensions/tasks/overview> [Standard]
71. A2A Protocol. *A New Chapter for A2A: Joining the Agentic AI Foundation*. A2A blog, 27 August 2026. <https://a2a-protocol.org/latest/blog/2026/08/27/a-new-chapter-for-a2a-joining-the-agentic-ai-foundation/> [Standard (governance)]
72. Microsoft. *What's new in Microsoft Entra Agent ID*. Product documentation, 1 May 2026 (updated 13 August 2026); general availability April 2026. <https://learn.microsoft.com/en-us/entra/agent-id/whats-new-agent-id> [Product release]
73. Microsoft. *Administrative relationships in Microsoft Entra Agent ID (Owners, sponsors, and managers)*. Product documentation, updated 3 September 2026. <https://learn.microsoft.com/en-us/entra/agent-id/agent-owners-sponsors-managers> [Framework guidance]

74. Microsoft (Nirav Shah, Rob Lefferts, and Jason Roszak). *Microsoft Agent 365, now generally available, expands capabilities and integrations*. Microsoft Security Blog, 1 May 2026. <https://www.microsoft.com/en-us/security/blog/2026/05/01/microsoft-agent-365-now-generally-available-expands-capabilities-and-integrations/> [Product release]
75. Google. *Agent Payments Protocol (AP2)*. Specification v0.2.0, 28 April 2026. <https://ap2-protocol.org/> [Standard (draft)]
76. Mastercard. *Verifiable Intent*. Open specification, draft v0.1 (February 2026). <https://github.com/agent-intent/verifiable-intent> [Standard (draft)]
77. OpenAI and Stripe. *Agentic Commerce Protocol: Delegate Payment API*. Specification version 2026-04-17 (beta). https://github.com/agentic-commerce-protocol/agentic-commerce-protocol/blob/main/spec/2026-04-17/openapi/openapi.delegate_payment.yaml [Standard (draft)]
78. Visa. *Visa Introduces Trusted Agent Protocol: An Ecosystem-Led Framework for AI Commerce*. Press release, 14 October 2025. <https://usa.visa.com/about-visa/newsroom/press-releases.releaseId.21716.html> [Standard (draft)]
79. FIDO Alliance. *FIDO Alliance to Develop Standards for Trusted AI Agent Interactions*. Press release, 28 April 2026. <https://fidoalliance.org/fido-alliance-to-develop-standards-for-trusted-ai-agent-interactions/> [Standard (programme)]
80. European Parliament and Council of the European Union. *Directive (EU) 2024/2853 of 23 October 2024 on liability for defective products*. Directive; applies to products placed on the market after 9 December 2026. <https://eur-lex.europa.eu/eli/dir/2024/2853/oj> [Regulatory text]
81. European Parliament. *Legislative Observatory: 2022/0303(COD) AI Liability Directive*. Procedure file; proposal withdrawn by Commission 6 October 2025. [https://oeil.secure.europarl.europa.eu/oeil/en/procedure-file?reference=2022/0303\(COD\)](https://oeil.secure.europarl.europa.eu/oeil/en/procedure-file?reference=2022/0303(COD)) [Regulatory text (procedure)]
82. Don Jergler. *Insurer Interest in AI Exclusions Growing as Risk Becomes Omnipresent*. Claims Journal, 20 July 2026 (secondary source for the ISO endorsements and Berkley exclusion). <https://www.claimsjournal.com/news/national/2026/07/20/338950.htm> [Survey (trade press)]
83. Armilla AI. *Affirmative AI Liability Insurance*. Product page, accessed September 2026. <https://www.armilla.ai/> [Product release]
84. Munich Re. *aiSure: Insurance for AI Performance Risk*. Product page, accessed September 2026. <https://www.munichre.com/en/solutions/for-industry-clients/insure-ai.html> [Product release]
85. HSB (Munich Re). *HSB AI Liability Insurance*. Product page, 2026. <https://www.munichre.com/hsb/en/products/ai-liability-insurance.html> [Product release]
86. Testudo. *AI Liability Insurance*. Product page, accessed September 2026. <https://www.testudo.co/> [Product release]
87. Artificial Intelligence Underwriting Company. *AIUC-1 Changelog: 15 July 2026 release*. Standard changelog. <https://standard.aiuc-1.com/changelog/> [Standard]
88. AIUC. *ElevenLabs Achieves AIUC-1 Certification*. Announcement, 18 February 2026. <https://www.aiuc-1.com/research/elevenlabs-achieves-aiuc-1-certification> [Product release]
89. K. J. Kevin Feng, David W. McDonald, and Amy X. Zhang. *Levels of Autonomy for AI Agents*. arXiv:2506.12469, June 2025 (revised July 2025); published in the Knight First Amendment Institute's *AI and Democratic Freedoms* essay series. <https://arxiv.org/abs/2506.12469> [Academic paper]

Annex Revision History

REVISION	DATE	CHANGES
2026-10	October 2026	First edition, published with VAOM v5.0. Carries the landscape and alignment content previously in whitepaper Section 13, updated and verified against primary sources as of 27 September 2026. New: GDPR Article 22 and explanation mapping; UK Articles 22A to 22D; US model risk guidance SR 26-2 and state decision-rights regimes; MAS SAFR, China, Korea, and OECD alignment; liability and insurance; agent payment protocols; MCP 2026-07-28, AIMS, and identity platform mappings. Corrected: Agent Control Standard version history; Entra Agent ID availability.

