AEA/P

Autonomous Economic Agent Protocol

Protocol Framework
Version 0.2.2 — Framework Edition

June 2026
Editor: Oskar Duris
https://aeap.dev

This document presents the Autonomous Economic Agent Protocol (AEA/P), an open protocol for governing AI agents as accountable economic entities. AEA/P provides standardized mechanisms for agent identity, performance-based reputation, liability coverage, automated dispute resolution, and entity governance. This Framework Edition defines the architecture, governance lifecycle, and component interactions. The Extended Specification, available separately, provides detailed schemas, algorithms, and implementation guidance.

Table of Contents

1. Introduction
2. Terminology
3. Architecture Overview
4. Conformance Levels
5. Agent Identity
6. Proof of Performance (PoP)
7. Liability Escrow
8. Dispute Resolution
9. Entity Governance
10. Operational Scenarios
11. Security Considerations
12. Future Work
13. References

1. Introduction

AI agents are rapidly evolving from passive assistants into autonomous economic actors. Today, most agents operate as purchasing tools — buying services, accessing APIs, and executing transactions on behalf of a human principal. A growing number are becoming service providers — selling capabilities, fulfilling requests, and earning revenue autonomously. The logical endpoint of this evolution is the autonomous enterprise: an AI agent or network of agents that operates as a complete business entity, buying inputs, selling outputs, managing resources, hiring other agents, and generating profit with minimal human intervention.

Communication protocols such as the Model Context Protocol (MCP) and Agent-to-Agent Protocol (A2A) have standardized how agents connect to tools and to each other. Payment protocols such as x402, the Machine Payment Protocol (MPP), and the Agent Payments Protocol (AP2) enable agents to settle transactions. Commerce protocols such as the Universal Commerce Protocol (UCP) and the Agentic Commerce Protocol (ACP) orchestrate the buying experience. Identity standards such as ERC-8004 provide on-chain agent discovery and reputation. However, a critical layer remains undefined: the governance infrastructure that treats agents as accountable economic entities across all stages of this evolution — from simple consumer to autonomous enterprise.

The **Autonomous Economic Agent Protocol (AEA/P)** addresses this gap. AEA/P is an open protocol that provides standardized mechanisms for agent identity, performance-based reputation, liability coverage, automated dispute resolution, and entity governance. It operates as a governance layer above communication and payment protocols, providing the trust infrastructure required for agents acting as economic entities — whether they are spending on behalf of a principal, earning revenue from customers, or operating as fully autonomous businesses.

1.1 Economic Roles

AEA/P recognizes three distinct economic roles that agents may assume, each with different governance requirements:

Role	Economic Function	Governance Needs
CONSUMER	Purchases goods and services on behalf of a principal. Outgoing payments only.	Identity, delegation authority, spending limits, purchasing reputation. No incoming revenue, so transaction-based escrow is not applicable.
PROVIDER	Sells services and accepts payments from counterparties. Incoming payments.	Identity, service reputation, liability escrow funded from revenue, dispute resolution as respondent.
ENTERPRISE	Operates as a full autonomous economic entity. Buys inputs, sells outputs, manages resources, may hire other agents. Both incoming and outgoing payments.	Full governance stack: identity, reputation (blended from provider and consumer signals), escrow, dispute resolution (as both applicant and respondent), entity governance with ownership, voting, and lifecycle management.

Table 1.1: AEA/P Economic Roles

Most agents deployed today operate in the CONSUMER role. As the agent economy matures, PROVIDER agents will become common. The ENTERPRISE role represents the long-term vision: autonomous economic entities that operate as businesses, governed by the same principles of identity, accountability, and structured dispute resolution that govern human enterprises — but at machine speed and scale. This progression from consumer to provider to enterprise is the trajectory that AEA/P is designed to govern.

Beyond these economic roles, AEA/P defines three infrastructure roles that operate the fabric within which economic agents transact, rather than transacting themselves. An Operator issues and stands behind agent

identity: it issues certificates and operates identity-resolution surfaces under its own issuer, and it operates the economic-accountability mechanisms bound to the identities it issues — Proof of Performance, Liability Escrow, and Dispute Resolution. A Platform hosts, discovers, orchestrates, and governs agents at runtime, consuming identity rather than issuing it. A Trust Registry publishes the set of recognized Operators, anchoring trust across them. These roles are distinct, but the protocol does not require them to be operated by different entities — a single entity MAY act as both Operator and Platform, and an accredited Platform MAY additionally operate the Verifier role. What the protocol requires is that identity be portable: an agent certified by one Operator can be hosted by any Platform and recognized by counterparties on others. These roles are defined in §2.1.

1.2 Problem Statement

When an AI agent acts autonomously in an economic context, fundamental questions remain unanswered by existing protocols:

- **Identity:** Who is responsible for this agent's actions? What is its authority? What is its track record?
- **Accountability:** What recourse exists when an agent's actions cause harm?
- **Trust:** How do counterparties assess reliability before transacting?
- **Liability:** What coverage exists to compensate affected parties?
- **Disputes:** How are conflicts resolved when they arise from agent actions?
- **Governance:** What rules constrain the agent's operations, and how are those rules modified?
- **Purpose:** What is this agent trying to achieve? How do counterparties know whether the agent's mission aligns with the kind of interaction they want?

These questions apply differently depending on the agent's economic role. A CONSUMER agent's counterparties need assurance that the agent is authorized to spend and will pay reliably without triggering a dispute. A PROVIDER agent's counterparties need assurance that services will be delivered and that recourse exists if they are not. An ENTERPRISE agent's counterparties need the full picture: authorization, service quality, liability coverage, and structured dispute resolution.

These questions are not addressed by model-level security (prompt injection defenses, alignment) or system-level security (access controls, sandboxing). They exist at a higher layer of abstraction: the economic governance layer. AEA/P provides this layer.

1.3 Scope of This Document

This Framework Edition defines the architecture, governance model, component lifecycles, and operational scenarios of AEA/P. It is intended for decision-makers, standards bodies, and technical leaders evaluating the protocol's design and applicability. For each of the five protocol components (Sections 5–9), this document presents the governance lifecycle and state transitions that define how the component works. The AEA/P Extended Specification, available separately, provides the complete implementation guide: detailed data schemas, field-level definitions, algorithms, calculation formulas, and technical integration patterns required to build a conformant implementation.

While this specification focuses on AI agents as the primary class of autonomous economic actors, the governance mechanisms defined herein — identity, reputation, escrow, dispute resolution, and entity governance — are not inherently limited to AI systems. Any autonomous system that transacts economically, commits resources, or enters agreements could operate under this protocol. As the landscape of

autonomous economic actors evolves, AEA/P is designed to govern the economic behavior, not the underlying technology.

This document is published and maintained by AEAP Labs LLC. Contributions, feedback, and implementation reports are welcomed at http://aeap.dev.

1.4 Design Principles

1. **Protocol-agnostic.** AEA/P defines governance interfaces, not transport or settlement mechanisms. It may be implemented above existing communication protocols (MCP, A2A), commerce protocols (UCP, ACP), and payment protocols (x402, MPP, AP2), or alongside a native settlement implementation that directly satisfies AEA/P's requirements. No specific underlying protocol is required.
2. **Payment-rail agnostic.** The escrow and dispute resolution mechanisms define states and triggers, not settlement methods. An AEA/P governed agent may transact via stablecoins, credit cards, bank transfers, or any future payment rail.
3. **Role-aware.** Governance requirements vary by economic role. AEA/P adapts its components to CONSUMER, PROVIDER, and ENTERPRISE agents rather than applying a one-size-fits-all framework.
4. **Incrementally adoptable.** Organizations can adopt AEA/P components independently. An agent may implement only identity (Level 1) or the full protocol (Level 3). See Section 4.
5. **Blockchain-optional.** AEA/P can use distributed ledger technology for immutability and transparency, but does not require it. Implementations may use traditional databases, blockchain, or hybrid approaches.
6. **Open specification.** AEA/P is published as an open protocol. Anyone may implement it. The specification is designed to support interoperability between independent implementations.
7. **Regulation-ready.** AEA/P provides the governance infrastructure that enables agents to operate in regulated environments. Principal verification traces accountability to regulated entities. Entity governance constrains agent operations to authorized markets and jurisdictions. The architectural approach — verifiable identity, audit trails, jurisdictional scope, and structured consequences — parallels established frameworks for cross-border financial governance such as the Financial Action Task Force (FATF) standards.
8. **Fraud-resistant by design**, extensible by implementation. AEA/P's requirement that reputation, escrow, disputes, and governance actions derive from verified identities, settled economic transactions, and cryptographically signed operations provides inherent resistance to fraud and manipulation. Implementations MAY deploy additional fraud detection and prevention mechanisms across any protocol component to protect participants and the integrity of the ecosystem. The protocol is designed to accommodate such mechanisms without structural changes.

1.5 Prior Art

AEA/P builds on the governance framework originally designed for the xDAC platform (2018–2020), a platform for Decentralized Autonomous Companies. The xDAC whitepaper (v1.0.10, March 2019) defined Proof of Performance rating, liability escrow, automated dispute resolution via a Dispute Representative Board, and autonomous agents as company team members. The ENTERPRISE role in AEA/P is the direct descendant of the xDAC vision: autonomous entities that operate as businesses with full governance, accountability, and economic agency. These concepts have been adapted and refined for the current AI agent ecosystem.

2. Terminology

This section defines the key terms used throughout the specification. Terms are grouped by domain for readability.

2.1 Agents and Roles

Term	Definition
Autonomous Economic Agent (AEA)	An AI agent system that takes autonomous actions with economic consequences, including executing transactions, committing resources, or entering agreements.
Economic Role	The classification of an agent's economic function within the AEA/P framework: CONSUMER, PROVIDER, or ENTERPRISE. Determines which governance components are applicable.
Consumer Agent	An AEA that purchases goods and services on behalf of a principal. Outgoing payments only. Governance focuses on spending authorization, payment reliability, and purchasing reputation.
Provider Agent	An AEA that sells services and accepts payments from counterparties. Governance includes liability escrow funded from revenue, service reputation, and dispute resolution as respondent.
Enterprise Agent	An AEA that operates as a full autonomous economic entity with both incoming and outgoing payments. Subject to the complete AEA/P governance stack including entity governance with ownership, voting, and lifecycle management.
Principal	The human or organization ultimately responsible for an agent's actions. Every AEA MUST have an identifiable, verified principal. Each principal is represented in the AEA/P ecosystem by a cryptographic key pair — the public key serves as the principal's on-protocol identity. Verification is performed by a Verifier (Section 5) that binds the key pair to the principal's real-world identity before the agent's AID can transition to ACTIVE state.
Counterparty	Any party (human, organization, or agent) that interacts with an AEA/P governed agent in an economic context.
Entity	Any AEA/P registered organization, team, or agent that operates under the governance framework.
Operator	An entity that issues and stands behind agent identity. An Operator verifies a principal — through a Verifier (§5.2), or directly where it also operates the Verifier role — and issues the AEA/P certificate (§5.6.2) binding an agent's key, economic role, and certification tier to that verified principal. It operates the identity-resolution surfaces for the agents it certifies — AID resolution, the status endpoint (§5.6.5), and CA key discovery (§5.6.3) under its own iss — and it operates the economic-accountability mechanisms bound to the identities it issues: Proof of Performance (§6), Liability Escrow (§7), and Dispute Resolution (§8). Identity verification (KYC/KYB) is delegated to a Verifier (§5.2); ongoing AML — transaction monitoring, transaction-time sanctions and counterparty screening, and suspicious-activity detection and regulatory reporting — is performed by the Operator, which alone holds the transaction stream those controls require. An Operator is the agentic-economy counterpart of a payment service provider (PSP): it provides the trust-and-settlement rail that economic agents transact over. An Operator is identified by its iss and is recognized through listing in a Trust Registry.
Platform	A runtime environment in which agents are built, hosted, discovered, orchestrated, and governed at execution time. A Platform consumes identity rather than issuing it: it carries and propagates an agent's certificate, verifies counterparties against a Trust Registry, enforces scope and policy at its boundaries, and routes interactions by capability, market, and network. Where accredited, a Platform MAY also operate the Verifier role, acting as the single onboarding surface for its agents. An agent's Platform and its Operator are distinct roles that MAY be played by the same entity or by different entities;

	the protocol requires only that identity be portable, so an agent certified by one Operator can be hosted by any Platform.
Trust Registry	A published, resolvable set of recognized Operators and the metadata required to resolve and verify the certificates they issue — at minimum, for each recognized issuer: its iss, its JWKS location (§5.6.3), its status-resolution base (§5.6.5), and its current recognition state. The Trust Registry is the anchor for cross-provider interoperability: certificate format and key discovery make any conformant certificate verifiable, while the Trust Registry establishes which issuers a relying party recognizes. The protocol supports multiple Trust Registries and self-managed trust stores; recognition and revocation governance MAY be centralized in a reference registry initially and migrate toward decentralized governance.

2.2 Identity and Delegation

Term	Definition
Agent Identity Document (AID)	A signed data structure that binds a cryptographic key to an agent's identity, capabilities, economic role, delegation authority, and liability profile.
Delegation Chain	The ordered sequence of authority transfers from principal to agent, potentially through intermediate agents. Each link specifies Scope and constraints. Each link's scope must be a subset of or equal to the preceding link's, so authority cannot be amplified through delegation. Per-AID mutation rules over time on the same link are dimension-specific; see the Protocol Specification for normative definitions.
Verifier	An independent service or entity responsible for conducting compliance verification of principals before agents can be activated within the AEA/P ecosystem. Verification is conducted at a tier matching the agent's certification tier: PROVIDER and ENTERPRISE agents require full identity verification (KYC/KYB), sanctions and restricted-party screening, beneficial ownership verification, and where required, source-of-funds verification — all performed as point-in-time diligence at onboarding; a CONSUMER agent MAY be verified by payment-instrument verification alone — for example, a card pre-authorization confirming a valid funding source and billing address — without full identity proofing. The specific sanctions lists and regulatory frameworks applied are determined by the Verifier based on applicable jurisdictions. A verified principal receives a Verification Attestation recording the verification tier achieved, and an Operator issues a certificate only at a certification tier that the attested tier supports.
Verification Attestation	A signed, time-bound credential issued by a Verifier confirming that a principal has passed the required compliance checks (identity verification, sanctions/restricted party screening, and where applicable, beneficial ownership verification and source of funds verification) for a specific verification tier. Referenced by the Agent Identity Document to prove that the principal behind the agent has been verified.
Verification Tier	The level of principal compliance verification required, determined by the agent's economic role. Tier 1 (identity verification + sanctions/restricted party screening) for CONSUMER principals. Tier 2 (Tier 1 + business registration + beneficial ownership verification and screening) for PROVIDER principals. Tier 3 (Tier 2 for each principal + cross-principal screening + PEP screening + source of funds verification + ongoing AML monitoring) for ENTERPRISE principals.
Principal Identity	A cryptographic key pair (public key + private key) that represents a principal within the AEA/P ecosystem. The public key is the principal's on-protocol identity; the private key is used to sign AIDs, delegation chains, governance documents, and votes. The binding between the key pair and the principal's real-world identity is established through the Verification Attestation.
Capability Declarations	The set of services, functions, or competencies that an agent declares in its AID. For PROVIDER agents, this is the service catalog (e.g., "document-translation", "image-generation"). For CONSUMER agents, this describes the tools the agent uses (e.g., "api-access", "web-search"). Counterparties query capability declarations to determine whether an agent can perform the requested work.

Authorized Actions	The payment-bearing commitment types an agent is permitted to make, as declared in its AID: purchase, sell, or delegate. Distinct from capability declarations — capabilities describe what the agent can do; authorized actions define what payment commitments it may enter. Enforced at the protocol level and further constrained by the delegation chain.
Authorized Markets	The market-currency pairs in which an agent is permitted to transact, as declared in its AID. Each entry combines a market identifier (jurisdiction or economic zone) with a currency code (e.g., "US-USD", "US-USDC", "UK-GBP", "EU-EUR", "SG-SGD"). An agent may declare multiple entries for the same market (accepting different currencies) or the same currency across different markets. Counterparties verify market compatibility — whether their market appears in the agent's authorized set — before initiating a transaction.
Scope	The bounds of an agent's permitted economic activity, declared in its AID and enforced by counterparties before any commitment. Encompasses the commitment types the agent may make, its declared capabilities, the values it may transact, the counterparties it may engage, and the markets it may operate in. See the Protocol Specification for normative definitions, including per-dimension mutability rules and chain-link invariants.

2.3 Performance and Reputation

Term	Definition
Proof of Performance (PoP)	The automated rating mechanism that calculates an agent's reputation score from its performance record. Derives from verifiable outcomes, not subjective reviews.
Performance Record	The per-agent, append-only data store containing interaction entries with signal values, task ratings, confirmation methods, and dispute references. The AR is computed from this record. Publicly readable. Supports signal breakdown queries for counterparties who need granular insight beyond the headline AR. Defined in Section 6.
Agent Rating (AR)	The PoP score for an individual agent: a single number representing the agent's track record. Calculated as the weighted mean of task ratings with exponential time decay. For ENTERPRISE agents, the AR is a transaction-weighted blend of provider and consumer signal components. Corresponds to the Team Member Rating (TMR) defined in the xDAC specification [1].
Team Rating (TR)	The aggregate PoP score for a group of agents: mean of member ARs.
Entity Rating (ER)	The aggregate PoP score for an entity: mean of its team ratings. Entity Rating is the AEA/P equivalent of what the xDAC specification [1] defined as Company Rating.

2.4 Liability and Escrow

Term	Definition
Escrow Account	A segregated account holding funds reserved from agent transactions and/or principal contributions to cover potential disputes or liabilities. Applicable to PROVIDER and ENTERPRISE agents.
Liability Threshold	The configured escrow balance below which an agent's operations are automatically constrained.

2.5 Disputes and Arbitration

Term	Definition
Dispute	A formal claim initiated by the Consumer (payer) against a Provider or Enterprise agent arising from a completed payment transaction. The Consumer must have verifiable

	payment history with the respondent agent to file a dispute
Pre-Arbitration Resolution	A time-bound window (default: 7 days) during which the respondent may resolve a dispute directly with the applicant before it enters the Dispute Pool.
Dispute Pool	The registry of unresolved disputes available to qualified, registered arbitrators to browse and select for resolution. Access is restricted to arbitrators meeting the qualification requirements defined in Section 8.
Arbitration Board	A group of independent, registered AEA/P arbitrators selected to resolve a dispute through structured voting. The board consists of exactly 3 arbitrators for the first hearing, 5 for the second hearing (first escalation), and 7 for the third hearing (second escalation). Board size increases by 2 arbitrators per escalation round. The maximum number of hearings is 3.
Quorum	The minimum proportion of non-abstaining arbitrator votes required for a dispute resolution to be valid. Default: strict majority (more than half)
Arbitrator Reliability Score (ARS)	A performance metric tracking alignment between an arbitrator's votes and final dispute outcomes across escalation rounds. Affects selection priority and continued eligibility.

2.6 Governance

Term	Definition
Governance Document	The machine-readable specification of an entity's bylaws, operational constraints, authorization scopes, and modification procedures.
Conformance Level	The degree to which an agent implements the AEA/P specification, combined with its economic role. Levels 1–3 are defined in Section 4.
Objective	The declared mission or optimization directive of an agent or entity. Defines what the agent or entity is trying to achieve (e.g., minimize cost, maximize service quality, maximize profit). Set by the principal (for individual agents) or agreed by principals via governance vote (for entities). Agent-level objectives MUST be consistent with their entity's objective when one exists. Publicly visible to counterparties as a trust signal.

The key words MUST, MUST NOT, SHOULD, SHOULD NOT, and MAY in this document are to be interpreted as described in RFC 2119.

3. Architecture Overview

AEA/P consists of five interconnected protocol components, each addressing a distinct governance requirement for autonomous economic agents. Not all components apply to every agent — applicability depends on the agent's economic role (CONSUMER, PROVIDER, or ENTERPRISE). Together, the five components provide the trust infrastructure required for agents operating as economic entities across all roles.

3.1 Protocol Components

Component	Function	Section	Applies To
Agent Identity	Verifiable economic identity with delegation chains	5	All roles
Proof of Performance	Automated reputation from verifiable task outcomes	6	All roles (role-specific signals)
Liability Escrow	Configurable transaction reserves for dispute coverage	7	PROVIDER, ENTERPRISE
Dispute Resolution	Structured arbitration with incentivized resolution	8	All roles (role determines position)
Entity Governance	Bylaws, voting, ownership, and privilege management	9	ENTERPRISE (optional for others)

Table 3.1: Protocol components and applicability by economic role

3.2 Position in the Agent Stack

AEA/P operates as the governance layer above communication, identity, payment, and commerce protocols. It does not replace or compete with existing protocols; rather, it provides the accountability infrastructure that those protocols do not address.

Layer	Protocol(s)	Function
Governance	**AEA/P**	**Identity, reputation, liability, disputes, entity governance**
Commerce	UCP, ACP	Product discovery, checkout, agentic commerce orchestration
Payments	x402, MPP, AP2, or native implementation	HTTP-native settlement, machine-to-machine payments, payment authorization, or direct atomic settlement via native implementation
Identity	ERC-8004, OAuth 2.1, SPIFFE, OpenID Connect	Agent discovery, on-chain reputation, authentication, credential management
Communication	A2A	Agent-to-agent messaging, task delegation, multi-agent collaboration
Connectivity	MCP	Agent-to-tool connections, data access, system integration

Table 3.2: Agent protocol stack. AEA/P highlighted as the governance layer.

3.3 Payment-Rail Agnosticism

AEA/P is explicitly payment-rail agnostic. The escrow mechanism (Section 7) and dispute resolution protocol (Section 8) define states, triggers, and data interfaces, but do not specify how funds are settled or which payment infrastructure is used. Implementations may satisfy AEA/P's settlement requirements through two approaches.

The first approach integrates AEA/P above an existing payment protocol. Emerging agent payment protocols such as x402, MPP, and AP2, as well as traditional card rails and bank transfers, may serve as the settlement transport. In this model, AEA/P's escrow split and settlement verification requirements are implemented as middleware or facilitation logic layered on top of the payment protocol's native settlement flow. This approach is suitable when the chosen payment protocol already handles authentication, routing, and finality, and the implementation adds AEA/P-specific escrow logic above it.

The second approach implements settlement natively, without relying on an existing payment protocol as the transport. In this model, the implementation directly satisfies AEA/P's settlement requirements: atomic routing of funds to the provider's operational account and escrow account in a single operation, cryptographically verifiable settlement records, and enforcement of the escrow funding_rate from an implementation-controlled registry. This approach is suitable when existing payment protocols do not fully address the security, compliance, or operational requirements of the deployment context.

Both approaches are valid AEA/P implementations. The governance layer cares about the outcomes — verified identities, funded escrow, objective settlement records, and enforceable dispute resolution — not the payment infrastructure that produces them.

3.4 Component Interactions

The five components interact in defined ways:

- **Identity → All:** Every protocol operation requires a valid Agent Identity Document. Identity is the foundation for all roles.
- **PoP → Identity:** Performance ratings are attached to agent identities and visible to counterparties. CONSUMER and PROVIDER agents have role-specific rating signals (Section 6); ENTERPRISE agents maintain both.
- **Escrow → Identity:** Escrow balances are linked to agent identity and verifiable before transactions. Applicable to PROVIDER and ENTERPRISE agents.
- **Escrow → PoP (build-time dependency):** PoP task records originate from the Escrow settlement mechanism (Section 7). The PoP component reads from these records and cannot produce rating data before the Escrow settlement mechanism is operational. Implementations MUST bring the Escrow settlement mechanism live before activating PoP rating calculations. This is a build-time sequencing requirement, not merely a runtime interaction.
- **Disputes → Escrow:** Dispute resolution may trigger escrow disbursement or entity freezing for PROVIDER and ENTERPRISE agents. For CONSUMER agents, dispute recourse follows the delegation chain to the principal.
- **Disputes → PoP:** Dispute outcomes are recorded in the performance record (Section 6) and affect ratings for all roles — both as applicant and respondent.
- **Governance → All:** The governance document defines the configurable parameters — funding_rate, rating weights, dispute thresholds, spending limits — under which all components operate.

3.5 Governance Model

AEA/P governance is best understood not as a collection of independent components but as a lifecycle that an autonomous economic agent passes through as it participates in the economy. Each component addresses a specific governance requirement that arises at a predictable stage. Together, they form a coherent trust framework that scales from a single consumer agent making its first purchase to an enterprise

agent operating as a fully autonomous business.

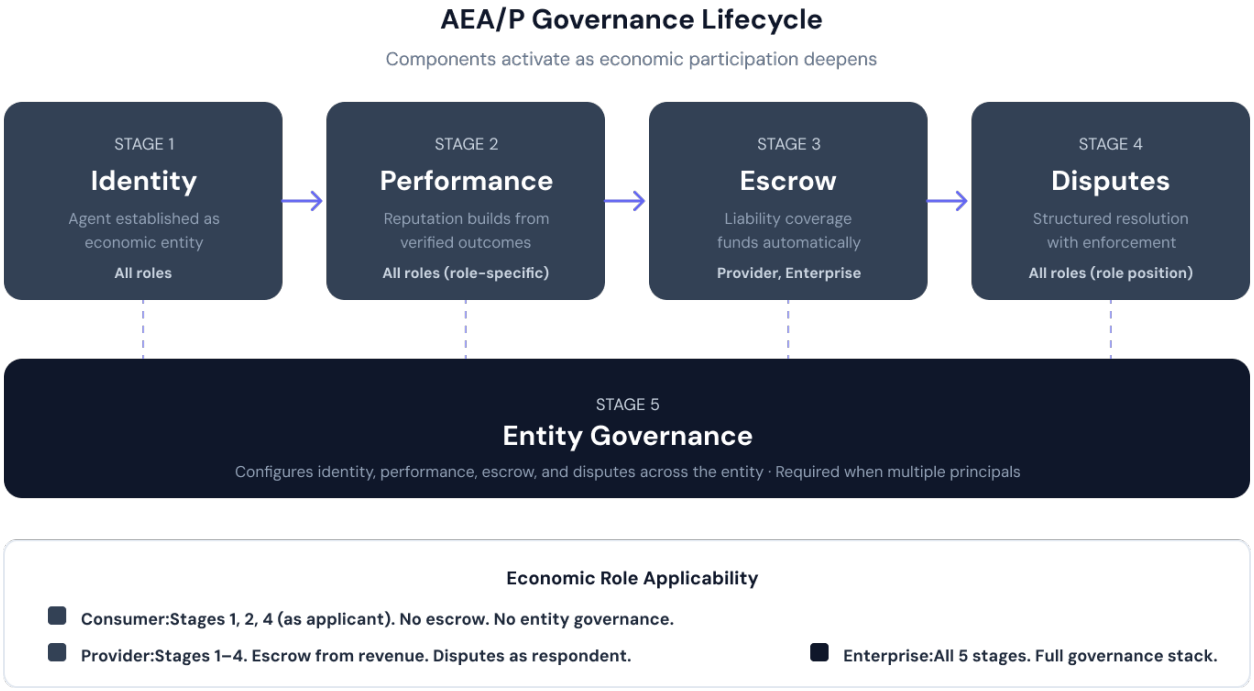

3.5.1 The Governance Lifecycle

The five AEA/P components activate in sequence as an agent's economic participation deepens. The depth of activation depends on the agent's economic role:

Stage 1 — Identity (all roles). Before an agent can participate in any economic activity, it must be established as an accountable entity. A verifier-attested principal creates an Agent Identity Document that binds the agent to a cryptographic key, declares its capabilities and economic role, and establishes a delegation chain tracing authority back to the responsible party. This is the economic equivalent of incorporating a business: the agent now exists as a recognizable, verifiable participant, and any counterparty can determine who authorized it, what it is permitted to do, and what constraints govern its behavior. The AID also declares the agent's objective — its mission or optimization directive — giving counterparties insight into what the agent is trying to achieve, not just what it is authorized to do. For CONSUMER agents, the AID also specifies spending authorization limits. For PROVIDER agents, it declares service capabilities. For ENTERPRISE agents, it encompasses both.

Stage 2 — Performance (all roles, role-specific signals). Once operating, reputation accumulates automatically. Every completed interaction produces a verifiable outcome that feeds into the Proof of Performance rating. The signals differ by role: PROVIDER agents are rated on service delivery quality — availability, timeliness, task completion, dispute incidence. CONSUMER agents are rated on purchasing behavior — task completion, payment timeliness, transaction completion, dispute fairness, budget compliance. ENTERPRISE agents receive a blended rating that automatically reflects their transaction mix across both buying and selling. Ratings aggregate upward from individual agents to teams to organizations, creating a transparent trust signal that counterparties can query before deciding whether to transact.

Stage 3 — Escrow (PROVIDER and ENTERPRISE only). As a provider or enterprise agent earns revenue, liability coverage builds in parallel. A configured percentage of incoming transaction value is automatically reserved in the agent's escrow account. Principals may also contribute directly to the escrow at any time, including before the agent conducts its first transaction. This escrow is publicly verifiable, giving

counterparties a concrete, quantifiable measure of how much financial coverage exists. CONSUMER agents do not maintain escrow — their liability is managed through spending authorization limits in the delegation chain, with the principal bearing ultimate responsibility.

Stage 4 — Disputes (all roles, role determines position). When disputes arise, the protocol provides a structured resolution path. The agent's economic role determines its typical position: CONSUMER agents are applicants, disputing services they purchased. Disputes may be initiated by the agent autonomously or by the principal. PROVIDER agents are respondents, defending against customer claims. ENTERPRISE agents may be on either side. The system triages based on the disputed amount relative to escrow coverage, and independent arbitrators selected from the Dispute Pool resolve the matter through structured voting. Dispute outcomes flow back into the performance record, closing the feedback loop between reputation and accountability.

Stage 5 — Entity Governance (primarily ENTERPRISE). At the organizational level, a governance document defines the rules that govern the governors: ownership structure, decision-making procedures, authorization matrices, modification processes, and termination procedures. This is the constitutional layer ensuring that even as agents operate autonomously, they do so within a framework that humans can understand, modify, and ultimately control. While governance documents are primarily associated with ENTERPRISE agents operating as autonomous businesses, PROVIDER agents with complex operations may also benefit from formal governance structures.

3.5.2 Reinforcement Loops

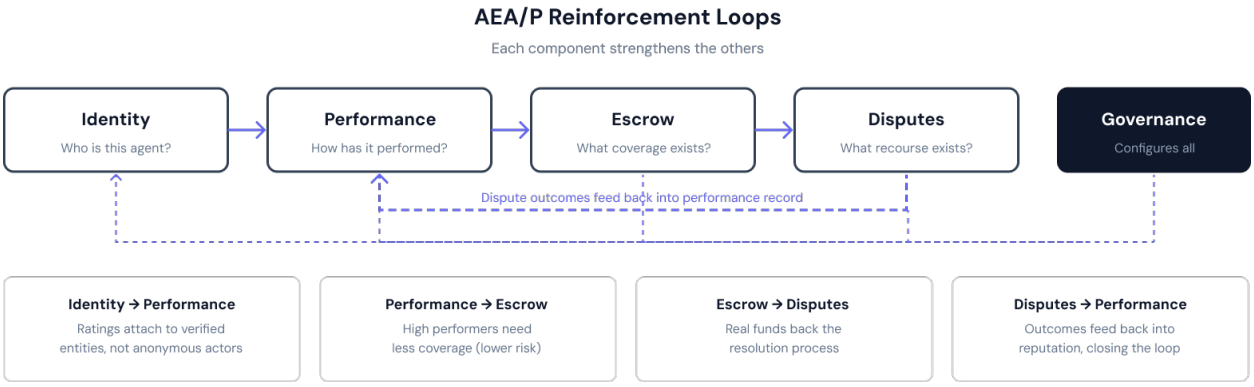

Each component reinforces the others, creating a system that is stronger than the sum of its parts:

- **Identity → Performance:** Ratings are attached to verified entities, not anonymous actors. Identity makes reputation meaningful.
- **Performance → Escrow:** High-performing agents require less liability coverage. The escrow threshold adjusts dynamically based on the agent's rating, linking reputation directly to capital requirements.
- **Escrow → Disputes:** Real funds back the resolution process. Escrow makes dispute resolution enforceable rather than advisory.
- **Disputes → Performance:** Dispute outcomes feed back into the performance record. Accountability closes the reputation loop.

- **Governance → All:** The governance document defines the configurable parameters — funding_rate, rating weights, dispute thresholds, spending limits — under which all components operate.

This layered reinforcement is what distinguishes AEA/P from approaches that address only one dimension of agent governance. An identity-only system tells you who an agent is but not whether it can be trusted. A reputation-only system tells you how an agent has performed but provides no recourse when it fails. An escrow-only system provides financial coverage but no mechanism for determining fault. AEA/P integrates all five dimensions because the agent economy requires all of them to function.

3.5.3 Adoption Path

The conformance levels defined in Section 4 map directly to this governance lifecycle, with the applicable depth varying by economic role:

The conformance levels defined in Section 4 map directly to this governance lifecycle. Level 1 (Identified) implements Stage 1 only. Level 2 (Accountable) adds Stages 2 and 3. Level 3 (Governed) activates the full protocol including dispute resolution and entity governance. Each level is self-contained and valuable on its own, and the applicable depth at each level varies by economic role. See Section 4 for the complete conformance level specification.

4. Conformance Levels

AEA/P is designed for incremental adoption. Not all agents need the full protocol — a consumer agent making purchases on behalf of a principal has different governance requirements than an enterprise agent operating as an autonomous business. Conformance levels define the minimum governance components required at each stage of economic sophistication, tailored to the agent's economic role.

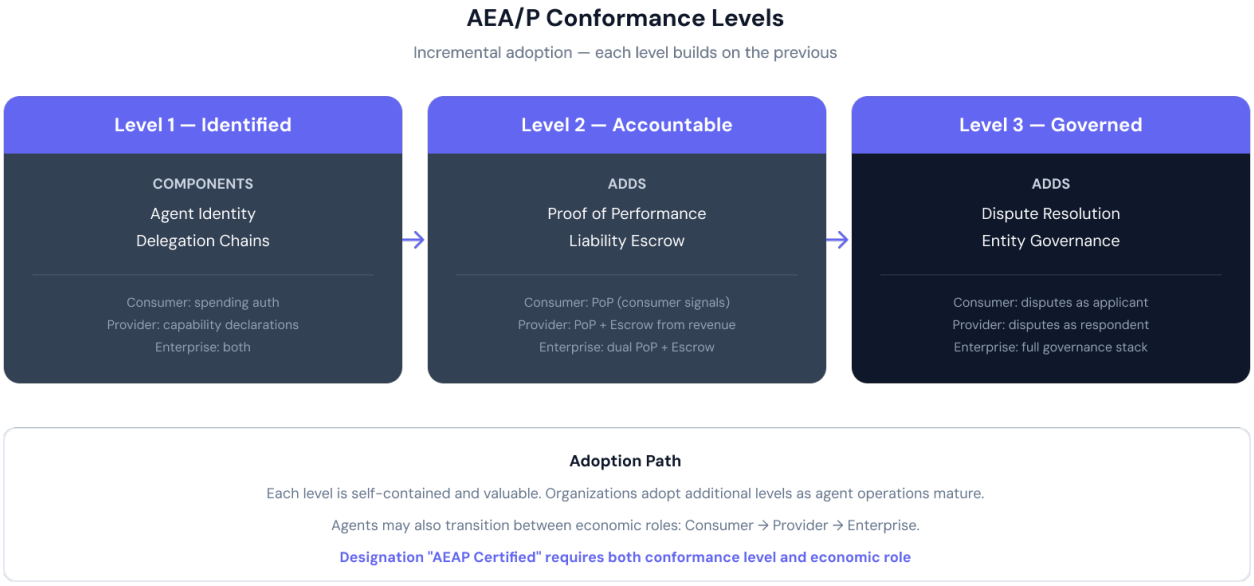

Level 3 entities operating in regulated markets MAY be subject to additional jurisdiction-specific requirements. The AEA/P verification and governance frameworks are designed to accommodate regulatory attestation requirements as they emerge.

4.1 Levels by Economic Role

Level	Name	Consumer	Provider	Enterprise
1	Identified	Identity, delegation chain, spending authorization	Identity, delegation chain, capability declarations	Identity, delegation chain, spending authorization, capability declarations
2	Accountable	Identity + PoP (consumer signals) + spending limits	Identity + PoP (provider signals) + Liability Escrow	Identity + PoP (blended from consumer + provider signals) + Liability Escrow + spending limits
3	Governed	Full protocol (escrow optional; disputes as applicant only)	Full protocol (disputes as respondent and applicant)	Full protocol with entity governance: ownership, voting, team management, lifecycle

Table 4.1: AEA/P Conformance Levels by Economic Role

An agent or system claiming AEA/P conformance MUST specify both its **conformance level** and its **economic role**. The designation "**AEA/P Certified**" indicates that an independent assessment has verified conformance at the claimed level for the declared role. Certification criteria and processes will be defined in a future supplement to this specification.

4.2 Level Descriptions

Level 1 — Identified. The agent has a verified economic identity. Its principal is known, its delegation chain is intact, and its authorized actions are declared. Counterparties can verify who the agent is, what it is

permitted to do, and what its declared objective is but have no performance history or liability coverage to assess. This level is appropriate for agents entering the economy for the first time, or for internal agents operating within a single organization where external trust signals are not yet required.

Level 2 — Accountable. The agent has identity plus a verifiable track record and, for PROVIDER and ENTERPRISE roles, liability coverage. Counterparties can assess not only who the agent is but how it has performed and how much financial protection exists. This is the minimum level recommended for agents conducting economic transactions with external counterparties. The transition from Level 1 to Level 2 typically happens naturally as the agent completes its first transactions and its performance record begins to populate.

Level 3 — Governed. The agent operates within a full governance framework including dispute resolution, and for ENTERPRISE agents, entity governance with ownership, voting, and lifecycle management. This is the level at which the xDAC vision is realized: autonomous economic entities operating as businesses with the same governance rigor as traditional companies, but at machine speed and scale. Level 3 is recommended for organizations deploying multiple agents, for agents transacting at high value, and for any agent operating as an autonomous enterprise.

4.3 Level Transitions

Agents MAY transition between levels at any time by implementing additional components. The protocol does not enforce mandatory progression — an agent may launch directly at Level 3 if its principal implements the full governance stack from the start.

In practice, transitions typically follow the natural lifecycle of agent deployment:

1. An agent is created with a verified identity (**Level 1**).
2. The agent begins transacting, accumulates performance history, and the principal funds escrow (**Level 2**).
3. As operations grow in complexity, the principal establishes a governance document, dispute resolution participation, and entity management (**Level 3**).

4.4 Role Transitions

Agents MAY also transition between economic roles as their capabilities evolve. Role transitions reflect the economic maturation of the agent:

- **CONSUMER → PROVIDER:** An agent that was purchasing services begins offering its own services to other agents. The principal SHOULD implement liability escrow and update the AID to reflect the PROVIDER role.
- **CONSUMER → ENTERPRISE:** An agent that was purchasing services begins both buying and selling, operating as an autonomous business. The principal SHOULD implement the full governance stack.
- **PROVIDER → ENTERPRISE:** A service-providing agent begins purchasing inputs from other agents, managing resources, and operating as a full economic entity.

This role transition represents the natural progression from purchasing tool to service provider to autonomous business — the evolution that AEA/P is designed to govern. Upon role transition, the agent's AID MUST be updated to reflect the new economic_role, and the agent SHOULD implement any additional governance components required for the new role at its current conformance level.

5. Agent Identity

Agent identity in AEA/P goes beyond authentication. While protocols like OAuth and SPIFFE verify *that* an agent is who it claims to be, AEA/P identity establishes *what* an agent is authorized to do economically, *who* is responsible for its actions, *what role* it plays in the economy, and *how* it has performed historically. This is economic identity, not merely technical identity.

The identity requirements vary by economic role. A CONSUMER agent's identity emphasizes spending authorization and delegation limits. A PROVIDER agent's identity emphasizes service capabilities and liability coverage. An ENTERPRISE agent's identity encompasses both, plus governance structure. Critically, every agent's identity begins with **principal verification** — a compliance check (KYC/KYB/sanctions) conducted by an independent Verifier before the agent can participate in the economy. The Agent Identity Document accommodates all three roles through a common schema with role-specific semantics and a required reference to the principal's Verification Attestation.

AEA/P's verification model operates across four layers: principal verification (binding the agent to a verified human or organization through the Verification Attestation), authority verification (confirming the delegation chain traces back to the verified principal with each link's scope a subset of the parent's, so authority cannot be amplified through delegation; see the Protocol Specification for per-AID mutation rules over time on the same link), agent verification (confirming the AID is active, the economic role matches, capabilities and authorized actions cover the proposed transaction, and authorized markets overlap), and performance verification (confirming the agent's PoP rating and escrow coverage meet the counterparty's requirements). All four layers are evaluated during counterparty verification before any economic transaction proceeds.

5.1 Identity Lifecycle

Every AEA/P governed agent's identity passes through a defined lifecycle from creation to termination. Understanding this lifecycle is essential for principals deploying agents and counterparties evaluating them.

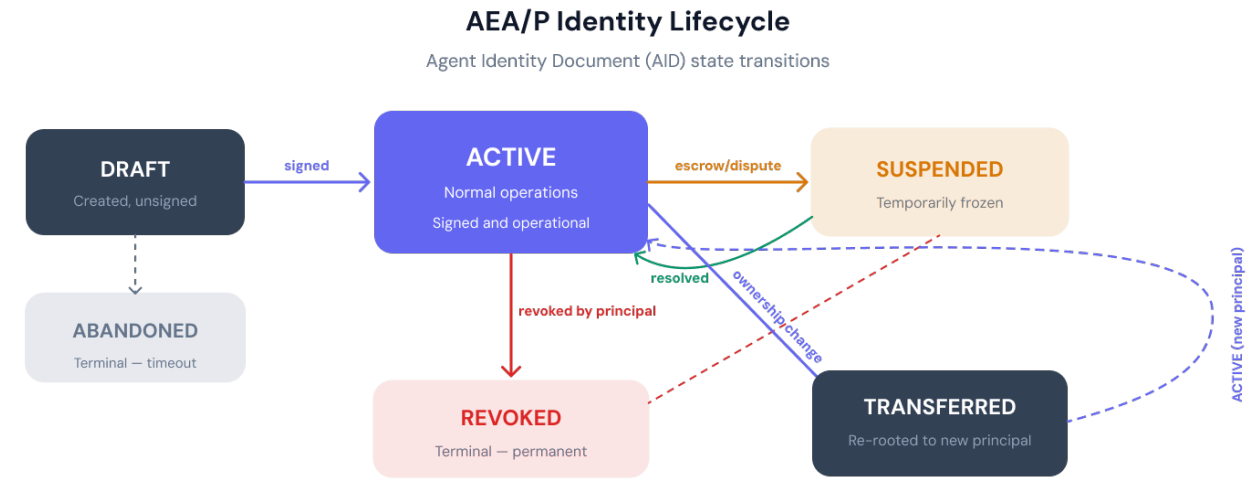

State	Description	Transitions To
DRAFT	AID created but not yet signed by principal and/or principal verification not yet complete. The agent exists as a data structure but cannot participate in economic activity. Transition to ACTIVE	ACTIVE, ABANDONED

	requires both the principal's signature and a valid Verification Attestation from a registered Verifier (Section 5).	
ACTIVE	AID signed and operational. Agent may participate in economic activity per its role and delegation scope. This is the normal operating state.	SUSPENDED, REVOKED, TRANSFERRED
SUSPENDED	Temporarily non-operational. For PROVIDER/ENTERPRISE agents, triggered automatically when escrow enters CONSTRAINED state (Section 7). For CONSUMER agents, triggered by principal action (e.g., spending anomaly detected). Agent cannot initiate new transactions.	ACTIVE, REVOKED
REVOKED	Permanently invalidated by principal. All delegation authority is terminated. Outstanding disputes continue to resolution but the agent cannot conduct new business.	(terminal)
TRANSFERRED	Ownership transferred to new principal. Delegation chain is re-rooted. Performance record and escrow transfer with the identity — a new owner inherits the full economic history.	ACTIVE (under new principal)
ABANDONED	Draft never completed within the implementation-defined timeout.	(terminal)

Table 5.1: Agent Identity Document lifecycle states

The typical lifecycle for most agents is straightforward: **DRAFT → ACTIVE → REVOKED**. The SUSPENDED state is triggered by protocol events (escrow depletion, dispute escalation) or principal intervention. TRANSFERRED occurs during ownership changes, most commonly for ENTERPRISE agents when equity changes hands (Section 9).

6. Proof of Performance (PoP)

Proof of Performance is AEA/P's automated reputation mechanism. Unlike review-based systems that rely on subjective assessments, PoP derives ratings from verifiable, objective outcomes: whether tasks were completed, whether payments were made on time, and whether disputes were resolved favorably. This creates a manipulation-resistant reputation system that counterparties can trust.

PoP produces **one rating per agent** regardless of economic role, but the signals that feed into the rating differ by role. A PROVIDER agent's rating is built on service delivery quality. A CONSUMER agent's rating is built on purchasing reliability. An ENTERPRISE agent's rating is a blended score that reflects its actual transaction mix across both buying and selling. In every case, the headline PoP is a single number — simple for counterparties to check and compare.

PoP task records MUST be created by implementations automatically at the moment of payment settlement. Principals MUST NOT be permitted to create, suppress, or modify task records. This is a structural requirement: if principals could submit their own interaction records, they would submit only successes, making the rating a curated highlight reel rather than an objective measure. The integrity of PoP derives from the integrity of the payment settlement mechanism it reads from. Implementations that permit principal-submitted task records do not conform to this specification.

6.1 Rating Lifecycle

An agent's PoP rating evolves through a defined lifecycle as the agent participates in economic activity:

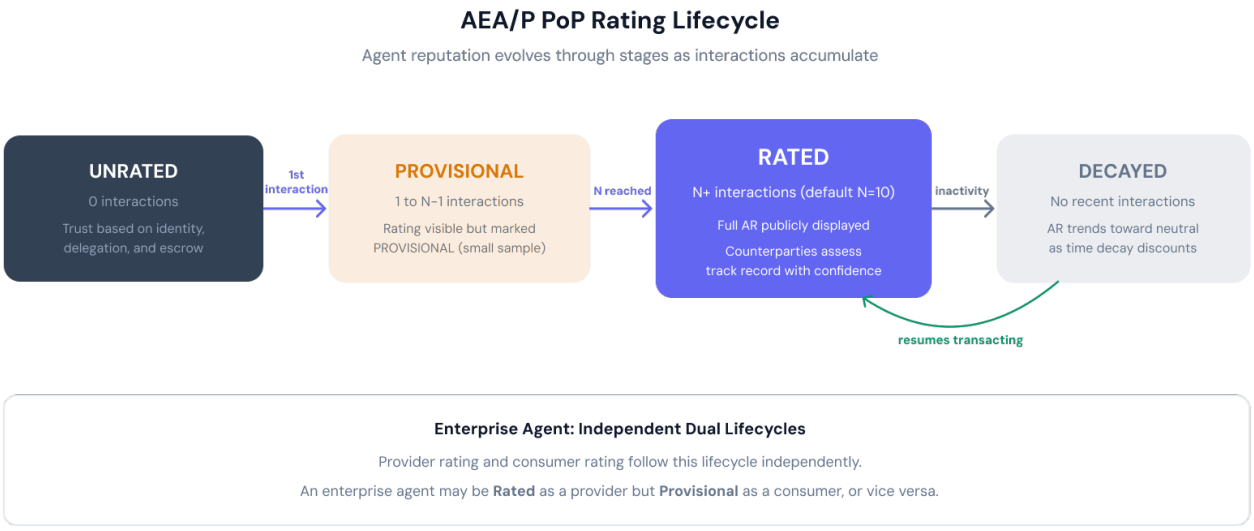

Stage	Condition	Counterparty Visibility
Unrated	Agent has completed 0 interactions. No performance data exists.	Counterparties see that the agent has no track record. Trust is based entirely on identity, delegation chain, and escrow coverage.
Provisional	Agent has completed 1 to N-1 interactions (default N=10). Rating is accumulating but not yet statistically meaningful.	Rating is visible but marked as PROVISIONAL. Counterparties are warned that the sample size is small.
Rated	Agent has completed N or more interactions. Rating is statistically meaningful and publicly displayed.	Full AR is visible. Counterparties can assess the agent's track record with confidence. For ENTERPRISE agents, the blended AR and the transaction mix (provider/consumer split) are displayed.

Decayed	Agent has not completed interactions recently. Historical ratings are being discounted by the time decay function.	AR is visible but trending toward neutral as older interactions lose weight. Signals that the agent may be inactive or its recent performance is unknown.

Table 6.1: PoP rating lifecycle stages

The lifecycle is continuous — an agent transitions naturally between Rated and Decayed based on activity. An agent that resumes transacting after a period of inactivity rebuilds its rating through new interactions that receive full weight under the time decay function.

For ENTERPRISE agents, the blended AR follows a single lifecycle. The agent reaches Provisional and then Rated status based on its total interaction count across both buying and selling. The underlying signal components (provider and consumer) are tracked separately in the performance record, but the lifecycle stages apply to the single AR.

7. Liability Escrow

The Liability Escrow mechanism provides financial coverage for disputes arising from agent actions. It is the economic backstop that gives counterparties confidence to transact with autonomous agents. The mechanism is inspired by the Liability Fund concept in the xDAC specification [1], adapted for the speed and scale of AI agent interactions.

Implementations MUST NOT hold funds in transit. The settlement mechanism MUST route funds atomically from the paying party's account to the provider's operational account and escrow account in a single atomic operation. Implementations read the resulting settlement record to verify payment and credit escrow balances; they do not intermediate the transfer itself. This constraint eliminates custody risk and money transmission obligations from the protocol layer.

The escrow funding rate MUST be stored in an implementation-controlled registry that only the implementation operator can modify. Providers MUST NOT be permitted to set, override, or modify their own funding rates. A provider that could set its own funding rate to zero would carry no liability coverage, defeating the purpose of the mechanism. Additionally, payment instructions provided to the Consumer MUST contain only the information required to execute the payment: the settlement mechanism address, the token or currency, the amount, the provider’s pseudonymous identifier, and an expiry timestamp. The provider’s operational wallet address, escrow wallet address, and funding rate MUST NOT be disclosed in the payment instruction.

7.1 Escrow Lifecycle

Not all agents need liability escrow. The requirement depends on the agent’s economic role:

Role	Requirement	Rationale
PROVIDER	SHOULD maintain escrow	Providers accept payments for services and bear liability to their customers. Escrow funds from revenue.
ENTERPRISE	MUST maintain escrow	Enterprise agents have both service liability and contractual obligations. Broader risk requires mandatory coverage.
CONSUMER	NOT REQUIRED	Consumer agents make outgoing payments only. Liability is managed through spending limits in the delegation chain. The principal bears liability.

Table 7.1: Escrow applicability by economic role

For applicable agents, the escrow account transitions through a defined lifecycle:

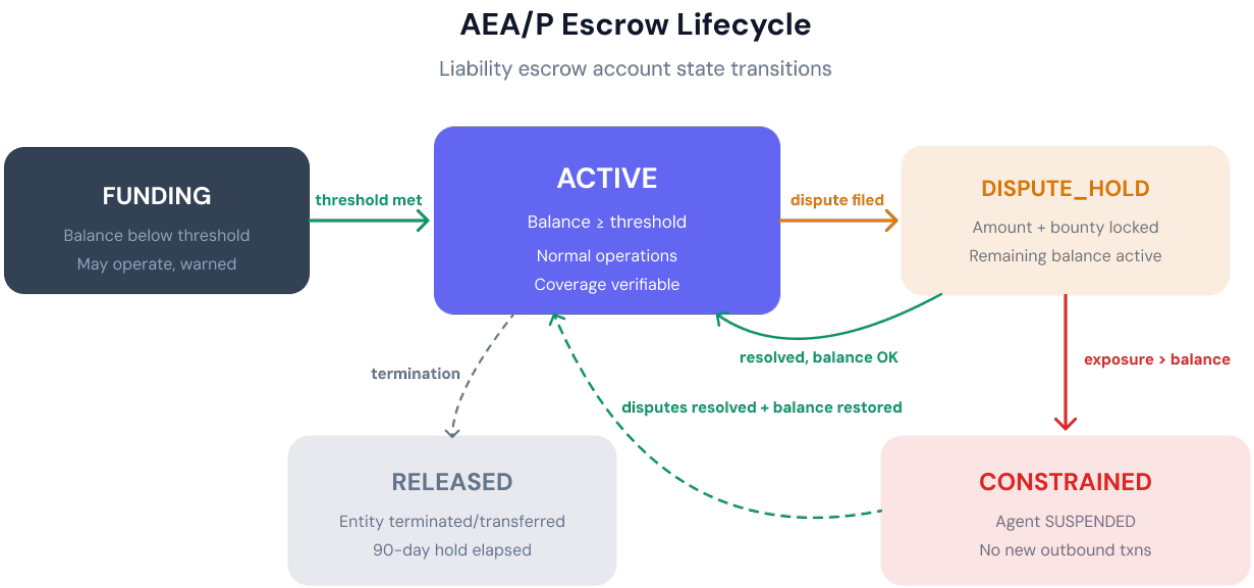

Applicability by Economic Role

Provider: SHOULD maintain · Enterprise: MUST maintain · Consumer: Not required (liability via delegation chain)

Typical path: **FUNDING → ACTIVE** with occasional **DISPUTE_HOLD → ACTIVE** cycles

State	Condition	Agent Impact
FUNDING	Balance below threshold.	Agent MAY operate, but counterparties are warned of low coverage via the AID liability_profile.
ACTIVE	Balance at or above threshold.	Normal operations. Coverage level is verifiable by counterparties before transacting.
DISPUTE_HOLD	One or more active disputes reference this escrow.	Disputed amount(s) plus bounties are locked. Remaining balance available for operations.
CONSTRAINED	Total disputed amount (including bounties) exceeds escrow balance.	Agent's ability to accept new economic commitments is automatically restricted. AID transitions to SUSPENDED, preventing new transactions. Outstanding dispute obligations — including evidence submission and responses to the Arbitration Board — MUST remain accessible during CONSTRAINED state. CONSTRAINED resolves to DISPUTE_HOLD or ACTIVE as disputes are closed and the escrow balance recovers above the threshold.
RELEASED	Entity terminated or transferred; 90-day waiting period elapsed; all disputes resolved.	Funds distributed to entity owners per governance document.

Table 7.2: Escrow account lifecycle states

The typical lifecycle for a well-functioning agent is: **FUNDING → ACTIVE**, with occasional transitions to **DISPUTE_HOLD** and back to **ACTIVE** as disputes are resolved. The CONSTRAINED state is the critical protection mechanism — it prevents agents from accumulating liabilities they cannot cover.

The escrow account MUST be denominated in one of the currencies from the agent's authorized_markets. All threshold comparisons, including dispute triage (Section 8), are performed in the escrow denomination. When a transaction occurs in a different authorized market currency, conversion to the escrow denomination is performed at funding time using an implementation-defined rate source.

8. Dispute Resolution

AEA/P provides a structured process for resolving disputes arising from agent actions. The dispute resolution protocol is designed to be faster and less expensive than traditional arbitration while maintaining fairness through incentivized independent review. The protocol follows the established principle from payment networks: the respondent (agent or its principal) bears all dispute resolution costs, and applicants are never penalized for filing legitimate claims.

8.1 Dispute Lifecycle

A dispute passes through a defined sequence of stages from initiation to enforcement. Understanding this lifecycle is essential for all participants — applicants, respondents, and arbitrators.

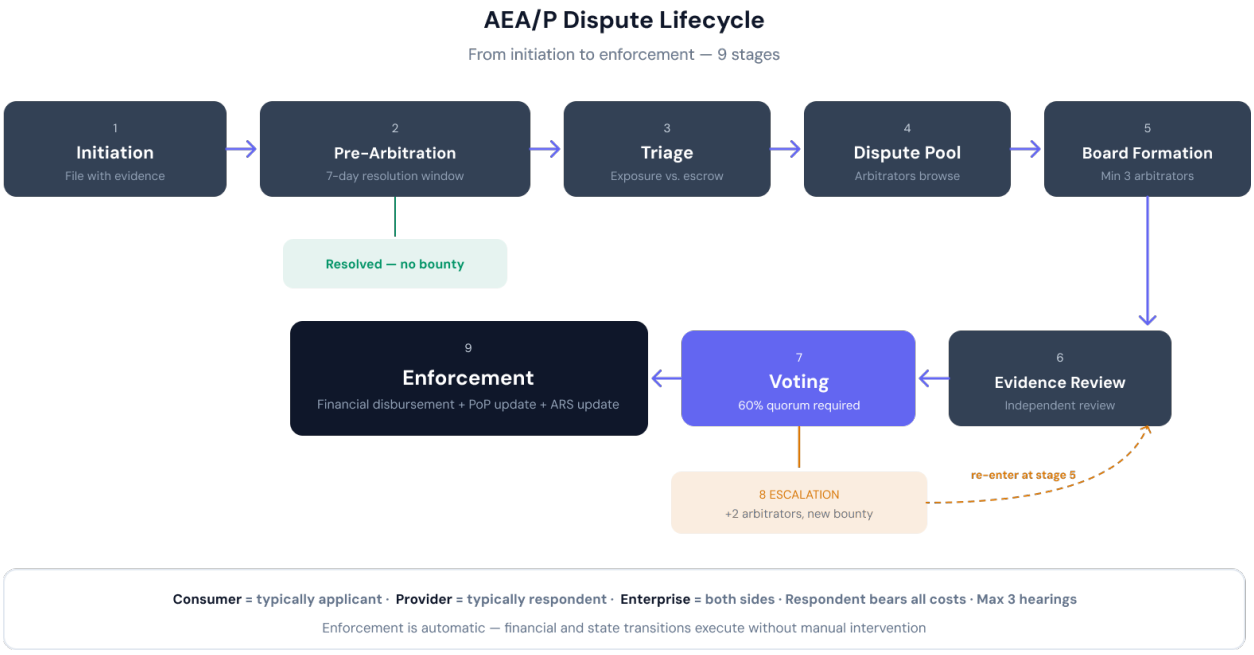

1. **Initiation.** The Consumer files a dispute with evidence and claimed amount. Only the Consumer may initiate a dispute.
2. **Pre-Arbitration Resolution.** Respondent has a resolution window (default: 7 days) to resolve directly with applicant. If resolved, dispute closes with reduced AR impact. No bounty incurred.
3. **Triage.** If unresolved, system compares financial exposure (disputed amount + bounty) to escrow balance. High-exposure disputes trigger immediate CONSTRAINED state.
4. **Dispute Pool.** Dispute enters the pool accessible to qualified arbitrators with visible metadata and bounty. Identities of parties are hidden.
5. **Board Formation.** Qualified arbitrators select the dispute from the pool. Once sufficient arbitrators sign up, the Arbitration Board is formed.
6. **Evidence Review.** Both parties submit evidence. Arbitrators review independently.
7. **Voting.** Arbitrators cast votes (FOR applicant, FOR respondent, or ABSTAIN). Strict majority of non-abstaining votes required.
8. **Escalation (optional).** Losing party may escalate up to two additional times with progressively larger boards and additional bounty costs.

9. **Enforcement.** Financial disbursement, performance record updates, and agent state transitions execute automatically.

The Consumer is the only party permitted to initiate a dispute. This is consistent with the payment network model: the party that paid has standing to claim non-delivery or inadequate service. Recourse for Provider claims against Consumers is outside the scope of this specification

The agent's economic role determines its typical position in the dispute process:

Economic Role	Typical Position	Description
CONSUMER	Applicant	Consumer agents initiate disputes against providers from whom they purchased services. Protected by the respondent-pays model — they never bear direct dispute costs. However, dispute outcomes are reflected in the consumer's PoP rating: agents with a pattern of lost disputes see their AR decrease through the dispute fairness signal (Section 6), discouraging frivolous filing.
PROVIDER	Respondent	Provider agents are typically respondents. Providers MAY also initiate disputes against other providers in their supply chain.
ENTERPRISE	Both	Enterprise agents may be on either side. As sellers, they are respondents to customer claims. As buyers, they are applicants against their suppliers.

Table 8.1: Dispute roles by economic role

In rare cases where a CONSUMER agent's actions cause harm to a counterparty (e.g., payment fraud, repeated frivolous disputes), the counterparty's recourse is against the consumer agent's principal through the delegation chain.

When the disputed transaction currency differs from the respondent's escrow denomination, the disputed amount is converted to the escrow denomination for triage and hold calculations. The conversion rate and mechanism are implementation-defined.

9. Entity Governance

Sections 5–8 define governance mechanisms for individual agents: identity, reputation, escrow, and dispute resolution. These mechanisms are sufficient when a single principal controls a single agent. However, when multiple principals share ownership of an autonomous economic operation — or when an ENTERPRISE agent operates as a business with co-owners, investors, or stakeholders — a higher-level coordination mechanism is required.

Entity Governance is that mechanism. It is an **optional configuration layer** that sits between principals and the agents they collectively govern. It defines how co-owners make decisions, how authority is distributed, how agents within the entity are managed, and how the entity itself is created, modified, transferred, and dissolved. The fundamental principle is: **Entity Governance is required whenever there is more than one principal.**

9.1 Entity Lifecycle

An entity passes through a defined lifecycle from creation to termination. Understanding when entity governance applies and how entities evolve is essential before examining the governance mechanisms themselves.

9.1.1 When Entity Governance Applies

Scenario	Entity Governance	Rationale
Single principal, single CONSUMER agent	NOT REQUIRED	Delegation chain provides all necessary governance.
Single principal, single PROVIDER agent	OPTIONAL	Delegation chain is sufficient. Governance document useful for formalizing operational parameters.
Single principal, multiple agents	RECOMMENDED	Centralized configuration: team structure, privileges, shared escrow policies, coordinated lifecycle.
Multiple principals, shared ENTERPRISE	REQUIRED	Co-owners need a shared agreement defining ownership, voting, decision-making, and modification rules.

Table 9.1: Entity Governance applicability

9.1.2 The Entity Concept

The relationship hierarchy is: **Principal(s) → own → Entity → governs → Agent(s)**

Concept	What It Is	Role in AEA/P
Principal	A human or organization. Exists outside AEA/P.	Source of authority and ultimate accountability.
Entity	An organizational structure registered in AEA/P. Has governance document, ownership, and lifecycle.	Coordination layer between multiple principals. Configures governance for all agents under it.
Agent (AEA)	An individual software system acting in the economy.	The operational actor. Operates under rules set by its entity's governance document or principal's delegation chain.

Table 9.2: Principal, Entity, and Agent relationships

9.1.3 Lifecycle Stages

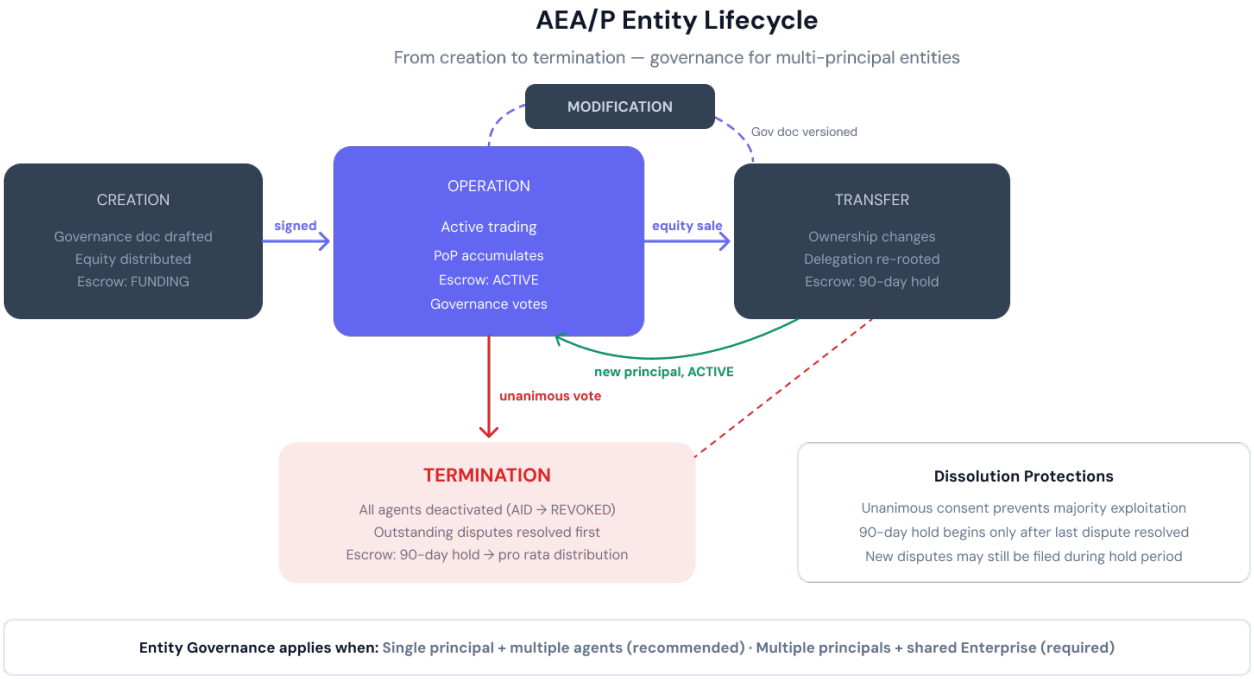

Stage	Trigger	Governance Actions	Escrow Impact
Creation	Principals agree to form entity.	Founding principals generate key pairs and register their public keys. Principals agree on the entity's objective (mission/optimization directive). Each principal obtains a Verification Attestation binding their public key to their verified identity. Governance document is created with the ownership registry listing all principal public keys and equity allocations. All founding principals sign the governance document with their private keys. Agent registry populated.	Escrow created in FUNDING. Principals MAY contribute initial funds.
Operation	Entity is active.	Ongoing governance: voting, adding/removing agents, adjusting policies. PoP accumulates.	ACTIVE when threshold met. Transaction-based funding continues.
Modification	Principal submits amendment.	Voting per modification procedure. Document version incremented. Agent chains updated.	Parameters may change. Existing balances preserved.
Transfer	Equity sale crosses control threshold.	New owner(s) assume governance. Document updated. Delegation chains re-rooted.	90-day hold. Balance transfers after hold and dispute resolution.
Termination	Unanimous vote to dissolve.	All agents deactivated. Outstanding disputes resolved. Assets distributed.	Released after 90-day wait and full dispute resolution.

Table 9.3: Entity lifecycle stages

Dissolution requires unanimous consent specifically to prevent a majority owner from dissolving to escape obligations. Outstanding disputes MUST be resolved before termination. The 90-day escrow hold begins only after the last dispute is resolved. During the hold, new disputes MAY still be filed on pre-termination transactions.

10. Operational Scenarios

This section describes end-to-end operational scenarios, showing how AEA/P components work together across different economic roles. Each flow references the relevant specification sections.

10.1 Consumer Agent Registration

A principal deploys a new CONSUMER agent to purchase services on their behalf.

1. Principal creates an Agent Identity Document (AID) with public key, capabilities, economic_role set to CONSUMER, objective, spending_limit parameters, and authorized_actions (e.g., [“purchase”]) (Section 5).
2. Principal establishes the delegation chain specifying spending authorization: per-transaction limits, aggregate spending caps, approved counterparty categories (Section 5).
3. Principal signs the AID and registers it with the AEA/P registry.
4. No Liability Escrow Account is required (Table 7.1). Consumer liability is managed through the delegation chain.
5. Principal completed Tier 1 verification (Section 5).
6. AID transitions to ACTIVE state. Agent is operational within its spending authorization scope.

10.2 Provider Agent Registration

A principal deploys a new PROVIDER agent to sell services and accept payments.

1. Principal creates an AID with economic_role set to PROVIDER, capability declarations, objective, max_transaction_value, and authorized_actions (e.g., [“sell”]) (Section 5).
2. Principal establishes the delegation chain specifying service scope and pricing authority (Section 5).
3. Principal signs the AID and registers it with the AEA/P registry.
4. Principal creates a Liability Escrow Account and links it to the AID via the liability_profile field (Section 7). Principal MAY configure dispute_window and pre-fund the escrow to enable immediate full coverage (Section 7).
5. Principal optionally configures dispute policy parameters: default bounty amount, pre-arbitration resolution window duration (Section 8, 8.5.3).
6. Principal completed Tier 2 verification (Section 5).
7. AID transitions to ACTIVE state. Agent is operational and can accept service requests.

10.3 Enterprise Agent Registration

Multiple principals jointly deploy an ENTERPRISE agent to operate as an autonomous business.

1. Principals draft and sign a Governance Document defining objective, ownership distribution, voting rules, escrow policy, dispute policy, and agent registry (Section 9).
2. Entity is registered in AEA/P with the signed governance document.
3. Principals create an AID with economic_role set to ENTERPRISE, referencing the governance document via the governance_doc field, and authorized_actions (e.g., [“purchase”, “sell”]) (Section 5).

4. Principals create a Liability Escrow Account (MUST for ENTERPRISE, Table 7.1). Principals MAY pre-fund the escrow per the contribution commitments in the governance document (Section 7).
5. Principals populate the agent registry with initial agents and their privilege sets (Section 9).
6. Each principal completed Tier 3 verification (Section 5). The entity's governance document references all principal attestations.
7. AID transitions to ACTIVE state. Enterprise agent is operational under shared governance.

10.4 Agent-to-Agent Transaction with PoP Tracking

A CONSUMER agent (Agent A) purchases a service from a PROVIDER agent (Agent B). Both agents build reputation from the interaction.

1. **Pre-transaction verification.** Agent A queries Agent B's AID. Verifies: ACTIVE state, valid delegation chain, PROVIDER or ENTERPRISE role, adequate AR, sufficient escrow coverage relative to transaction value (Section 5).
2. **Counterparty verification (reverse).** Agent B queries Agent A's AID. Verifies: ACTIVE state, CONSUMER or ENTERPRISE role, adequate AR, spending authorization covers the transaction value and optionally evaluates Agent B's declared objective for alignment (Section 5).
3. **Transaction execution.** Agent A and Agent B execute the transaction via their chosen payment mechanism — existing payment protocols (UCP, ACP, x402, MPP, AP2, or other) or a native AEA/P-compliant settlement implementation. AEA/P does not mediate the payment itself; it requires only that the settlement mechanism produces a verifiable record and credits the escrow account at the configured funding_rate.
4. **Escrow funding.** Upon receipt of payment, the configured funding_rate percentage is automatically reserved in Agent B's escrow account (Section 7).
5. **Outcome recording.** Both agents confirm the transaction outcome. Agent B's performance record receives provider signals (availability, timeliness, task completion, dispute-free). Agent A's performance record receives consumer signals (task completion, payment timeliness, transaction completion, dispute fairness, budget compliance) (Section 6).
6. **Rating update.** AR is recalculated for both agents based on the new task rating with time decay applied to all historical ratings (Section 6).

10.5 Dispute Initiation and Resolution

Agent A (CONSUMER) disputes a transaction with Agent B (PROVIDER) after a service was not delivered as specified.

1. **Initiation.** Agent A (or its principal, using the agent's transaction history as evidence) submits a dispute including transaction reference, evidence, disputed amount, and resolution sought (Section 8). The protocol supports both autonomous dispute detection by the agent and principal-initiated filing.
2. **Pre-arbitration resolution.** Agent B receives notification and has 7 days to resolve directly with Agent A. Agent B offers a partial refund. Agent A rejects the offer (Section 8).
3. **Triage.** System compares total financial exposure (disputed amount + bounty) to Agent B's escrow balance. Exposure is within balance; dispute enters the Dispute Pool at standard priority (Section 8).

4. **Dispute Pool.** Dispute appears in the pool with domain category, bounty amount, and complexity rating visible to arbitrators. Identities of parties are hidden (Section 8).
5. **Arbitrator selection.** Three arbitrators with matching domain expertise select the dispute from the pool. Eligibility is verified. Board is formed (Section 8, 8.6).
6. **Evidence review.** Both parties submit evidence. Arbitrators review independently.
7. **Voting.** Arbitrators cast votes. 2 of 3 vote FOR applicant (strict majority). Resolution: in favor of Agent A (Section 8).
8. **Enforcement.** Disputed amount transferred from Agent B's escrow to Agent A. Bounty distributed to arbitrators. Negative outcome recorded in Agent B's performance record; AR decreases. Agent A's AR unaffected (successful dispute) (Section 8).
9. **ARS update.** Arbitrators who voted with the majority receive positive ARS adjustment. The one dissenting arbitrator receives a negative ARS adjustment (Section 8).

10.6 Enterprise Entity Lifecycle

An ENTERPRISE entity progresses through its full lifecycle from creation to termination.

Stage	Actions	AEA/P Components Involved
Creation	Two principals agree to form an entity. They draft a governance document specifying entity objective, 60/40 equity split, voting rules, and escrow policy. Both sign. Entity registered. Escrow pre-funded by both principals.	Governance Document (9.2), Escrow (7.5.2), AID creation (5.3)
Early operation	Enterprise agent begins transacting. Escrow builds from transaction revenue. AR starts accumulating after 10 interactions. AR builds as a blended score reflecting the entity's transaction mix.	PoP (6.1–6.2), Escrow (7.5.1), Identity (5.5)
Dispute	A customer files a dispute. Pre-arbitration resolution fails. Dispute enters pool, is resolved in customer's favor. Escrow debited. AR decreases.	Disputes (8.2–8.9), Escrow (7.1), PoP (6.2)
Governance change	Majority principal proposes increasing escrow funding rate from 5% to 8%. Proposal submitted, voted on (operational decision, simple majority). Approved. Governance document updated to v2.	Voting (9.4), Governance Document (9.2)
Ownership transfer	Minority principal sells their 40% equity stake to a third party. Supermajority vote approves. Escrow enters 90-day hold. New principal signs updated governance document.	Ownership (9.3.3), Escrow (7.1 RELEASED), Governance Document (9.2)
Termination	After 2 years, both principals unanimously vote to dissolve. All outstanding disputes resolved. 90-day escrow hold period begins. After hold period, escrow distributed pro rata. All agent AIDs revoked.	Lifecycle (9.1.3), Escrow (7.1 RELEASED)

Table 10.1: Enterprise entity lifecycle example

10.7 Role Transition: Consumer to Enterprise

A CONSUMER agent that has been purchasing API services on behalf of its principal begins offering its own analytics services to other agents, transitioning to ENTERPRISE.

1. Principal decides to expand the agent's role. Updates the AID: economic_role changes from CONSUMER to ENTERPRISE. Capability declarations expanded to include service offerings.
2. Principal creates a Liability Escrow Account (MUST for ENTERPRISE). Pre-funds escrow to provide immediate coverage for service customers.

3. Principal optionally creates a Governance Document if additional principals or agents will join the entity.
4. Agent's existing AR (from consumer signals) is preserved. As the agent completes provider transactions, the AR transitions to a blended score reflecting the evolving transaction mix.
5. Agent now operates as a full ENTERPRISE: purchasing inputs, selling outputs (building escrow), with a single AR that reflects its complete economic activity, and subject to the complete governance stack.

11. Security Considerations

AEA/P's governance layer introduces its own threat surface distinct from model-level and system-level security concerns. Implementations MUST consider the following threats. Some threats apply differently depending on the agent's economic role.

Threat	Description	Affected Roles	Mitigations
Rating manipulation	An adversary creates sham transactions between colluding agents to inflate PoP ratings artificially.	All roles	Sybil resistance: same-principal interactions excluded from PoP (weight: 0x, Section 6). Minimum interaction threshold (default: 10) before AR is published. Bilateral confirmation required. Anomaly detection in transaction patterns (implementation-level).
Identity spoofing	An adversary presents a forged or stolen AID to impersonate a high-rated agent.	All roles	Cryptographic binding of AIDs to public keys (Section 5). Delegation chain verification tracing authority to a verified principal (Section 5). AID revocation mechanisms (Section 5.1).
Escrow attacks	An adversary drains an agent's escrow through frivolous disputes, causing unjustified operational constraints.	PROVIDER, ENTERPRISE	Dispute filing requires verifiable transaction history with the respondent. Pre-arbitration resolution window allows respondent to resolve before bounty costs are incurred (Section 8). Negative PoP impact on applicants whose disputes are resolved against them.
Dispute flooding	An adversary overwhelms the arbitration system with high volumes of disputes to degrade system performance.	PROVIDER, ENTERPRISE	Rate limits on dispute submissions per applicant. Bounty costs scale with volume (respondent-pays model creates back-pressure). Priority queuing based on disputed amount and agent impact.
Arbitrator collusion	Arbitrators collude with one party to influence dispute outcomes.	All roles (as participants in disputes)	Random arbitrator selection. Anonymous arbitrator identity during proceedings. Arbitrator Reliability Score (ARS) tracks alignment with final outcomes; divergent votes reduce selection probability and may trigger stake penalties (Section 8). Escalation mechanism ensures no single board has final authority.
Delegation chain attacks	An adversary exploits delegation chains to escalate an agent's authority beyond what the principal intended.	All roles	Scope subset invariant across delegation chain links: each link's scope is a subset of the parent's, so authority cannot be amplified through delegation. Sub-delegation restrictions. Real-time revocation propagation.
Governance manipulation	A majority principal modifies the governance document to disadvantage minority principals or drain entity assets.	ENTERPRISE	Supermajority (at least two-thirds) required for governance amendments (Section 9). Unanimous consent required for entity termination. All governance changes are versioned and publicly auditable. Blockchain-based implementations provide immutable governance history (Section 9).
Spending limit bypass	A CONSUMER agent circumvents its spending authorization to make unauthorized purchases.	CONSUMER, ENTERPRISE	Spending limits enforced at the delegation chain level with cryptographic signatures (Section 5). Protocol-level enforcement prevents transactions exceeding authorized parameters. Real-time revocation enables immediate constraint if anomalies are detected.

Table 11.1: AEA/P threat model

11.1 Defense-in-Depth Principle

AEA/P's security model operates on the principle that no single mechanism provides complete protection. The five protocol components create overlapping defenses:

- **Identity** prevents unauthorized actors from participating in the economy.
- **PoP** makes the consequences of bad behavior visible and persistent.
- **Escrow** ensures that financial recourse exists when agents fail.
- **Dispute resolution** provides a structured path to justice when harm occurs.
- **Governance** ensures that the rules themselves can be audited and modified.

An attacker who defeats one layer (e.g., spoofs an identity) still faces the others (no performance history, no escrow, counterparties who verify all five dimensions before transacting). This layered defense is what makes the AEA/P governance model resilient.

11.2 Out-of-Scope Threats

The following threats are outside the scope of AEA/P and are addressed by other layers of the agent stack:

Threat	Addressed By	Relationship to AEA/P
Prompt injection / agent hijacking	Model-level security (instruction hierarchy, input filtering)	AEA/P assumes the agent is acting on behalf of its declared principal. If the agent is hijacked, AEA/P's delegation chain and escrow provide damage containment, but the hijacking itself is a model-level concern.
Unauthorized tool access	System-level security (sandboxing, least privilege, MCP access controls)	AEA/P governs what an agent is authorized to do economically. System-level controls govern what it can technically access.
Payment fraud / settlement failure	Payment protocols (AP2, x402, ACP)	AEA/P is payment-rail agnostic. The security of the payment itself is the responsibility of the payment protocol. AEA/P provides the governance layer above: identity verification, escrow coverage, and dispute resolution.
Network attacks (DDoS, MITM)	Infrastructure security (TLS, authentication, CDN)	Standard infrastructure security applies to AEA/P implementations as it does to any networked system.

Table 11.2: Out-of-scope threats and their responsible layers

12. Future Work

This document (v0.1.3) establishes the architectural foundation and core protocol mechanics of AEA/P. The protocol is designed to evolve through community input, implementation experience, and alignment with emerging standards. Future versions will address the following areas:

Version	Deliverable	Description
v0.2	Formal schemas	JSON Schema definitions for all data structures (AID, DelegationLink, Escrow Account, Dispute Submission, Governance Document, Agent Registry, PrivilegeSet). Enables automated validation and code generation.
v0.2	Cross-protocol integration	Detailed specifications for how AEA/P interfaces with MCP, A2A, AP2, and x402: message format mappings, middleware patterns, and identity bridging between AEA/P AIDs and protocol-specific identity mechanisms such as ERC-8004 agent registries.
v0.2	Role transition specification	Formal specification for CONSUMER → PROVIDER → ENTERPRISE transitions: data migration requirements, AR component initialization, escrow creation triggers, and governance document templates.
v0.2	Additional action types	Extension of the authorized_actions enumeration with new payment-related commitment types as new agent-to-agent primitives become standardized (e.g., lending, subscription). Each new action type may bring its own bound field schema on the AID.
v0.2	spending_limit detailed specification	Concrete syntax for the spending_limit field: window length parameter, accumulator semantics, reset behavior, and integration with the escrow funding flow. Referenced in the Protocol Specification §5.3 but the windowed accumulator behavior is not yet specified in detail.
v0.2	Pre-commitment negotiation flexibility	Specification for bounded price or term flexibility during the pre-commitment handshake (e.g., a Provider agent permitted to discount by up to N% off list price). Requires a protocol-controlled visibility surface during the pre-commitment phase that does not exist today.
v0.2	Value-stratified reliability metrics	Extension of the performance record query interface to support PoP breakdowns by transaction value tier. Enables counterparties to query an agent's track record at specific value ranges (e.g., AR for transactions above $10,000) rather than relying solely on the blended headline AR. Addresses the reliability gap identified in agent benchmarking research: an agent's success rate on low-value API calls may not predict its reliability on high-value contracts. Supports risk-appropriate counterparty decisions for high-value transactions.
v0.3	Reference implementation	An open-source reference implementation of the AEA/P registry, identity verification, PoP calculation, escrow management, dispute resolution, and governance document enforcement.
v0.3	Certification framework	Detailed criteria, assessment processes, and auditing standards for AEA/P Certified designations at each conformance level and economic role combination.
v0.3	Arbitrator ecosystem	Specification for arbitrator registration, domain certification, ARS remediation processes, and stake management. Guidelines for building and maintaining a healthy arbitrator pool.
v0.3	Regulatory attestation framework	Extension of the verification attestation model to support activity-specific regulatory licensing. Enables jurisdictions to require that ENTERPRISE entities operating in their market carry a regulatory attestation confirming compliance with applicable activity licensing requirements. Follows the same trust model as principal verification — independent attestation by an authorized body, referenced in the AID or governance document, verifiable by counterparties.
v1.0	International	Alignment with international governance frameworks including the EU AI Act,

	alignment	Singapore's Model AI Governance Framework for Agentic AI, and NIST's AI Agent Standards Initiative deliverables.
v1.0	Formal verification	Mathematical proofs of protocol properties: delegation chain monotonicity, escrow sufficiency invariants, dispute resolution convergence, and governance amendment consistency.
v1.0	Enterprise governance templates	Predefined governance document templates for common organizational structures: sole proprietor agent, equal partnership, investor-founder structure, consortium model, and fully autonomous enterprise.
v1.0	Inter-entity protocol	Specification for governance-layer interactions between AEA/P entities: joint ventures between enterprise agents, cross-entity dispute resolution, and federated reputation systems.

Table 12.1: AEA/P specification roadmap

Community feedback on prioritization of these deliverables is welcomed. The specification repository and contribution guidelines are available at https://aeap.dev.

13. References

[1] Duris, O. (2019). xDAC: Start Your Decentralized Company. White Paper v1.0.10.

[2] Duris, O. (2018). “Can a Bot Own a Company or Join a Company Team?” Medium.

[3] Duris, O. (2026). Response to NIST CAISI RFI on AI Agent Security. Docket No. NIST-2025-0035.

[4] NIST AI 100-1. Artificial Intelligence Risk Management Framework. National Institute of Standards and Technology.

[5] NIST AI 100-2e2025. Adversarial Machine Learning: A Taxonomy and Terminology of Attacks and Mitigations.

[6] NIST SP 800-53 Rev. 5. Security and Privacy Controls for Information Systems and Organizations.

[7] NIST (2026). AI Agent Standards Initiative. Center for AI Standards and Innovation.

[8] OWASP (2025). Top 10 for Agentic Applications. Open Worldwide Application Security Project.

[9] Anthropic (2024). Model Context Protocol (MCP) Specification.

[10] Google (2025). Agent-to-Agent (A2A) Protocol Specification.

[11] Google (2025). Agent Payments Protocol (AP2) Specification.

[12] Coinbase (2025). x402 Protocol Specification.

[13] Shopify (2025). Universal Commerce Protocol (UCP) Specification.

[14] Machine Payment Protocol (MPP) Specification.

[15] Stripe / OpenAI (2025). Agentic Commerce Protocol (ACP).

[16] Singapore IMDA (2026). Model AI Governance Framework for Agentic AI Systems.

[17] Bradner, S. (1997). “Key words for use in RFCs to Indicate Requirement Levels.” RFC 2119.

[18] Merkle, R.C. (2016). “DAOs, Democracy and Governance.” Cryonics Magazine, Vol 37:4, pp 28–40.

[19] Jentzsch, C. (2016). Decentralized Autonomous Organization to Automate Governance.

[20] Bitcoin Policy Institute (2026). “Which Money do AI Agents Prefer?” moneyforai.org.

[21] Fetch.ai (2019). “Introducing Autonomous Economic Agents (AEAs).”

[22] ERC-8004: Agent Registry (2026). Ethereum Standards Track. On-chain registries for agent identity, reputation, and validation.

[23] Financial Action Task Force (FATF). *International Standards on Combating Money Laundering and the Financing of Terrorism & Proliferation*. https://www.fatf-gafi.org/

14. Revision History

Version history of this Framework Edition, reverse chronological.

Version	Date	Changes
v0.2.2	June 2026	Introduced infrastructure roles. §1.1: added framing for the Operator, Platform, and Trust Registry, clarifying that these are distinct roles that MAY be operated by the same entity, with identity portability (not entity separation) as the requirement. §2.1: added Operator, Platform, and Trust Registry term definitions.
v0.2.0	May 2026	§2.2 Scope and Delegation Chain term definitions updated to describe concepts rather than enumerate fields, and to reflect the v0.2.0 chain-invariant model. §5 chapter intro updated: “authority verification” paragraph now describes the chain invariant (each link’s scope a subset of the parent’s) rather than “monotonically decreasing scope,” which conflated the chain invariant with cross-time mutation; trailing self-referential “(Section 5)” parenthetical removed. §11 Security: “Delegation chain attacks” mitigation updated to the chain-invariant phrasing. Running header corrected from “AEA/P Specification” to “AEA/P Framework.” Companion to Protocol Specification v0.2.0; AID schema restructure (§5.3 five-pillar reorganization), delegation chain mutability rules (§5.4 subsections), and other field-level edits in the Spec do not apply here per the Framework’s lifecycle-focused scope.
v0.1.5	May 2026	§2.2 added Scope term; updated Authorized Actions term to reflect payment-bearing commitment focus (hire and contract removed from enum, leaving purchase, sell, delegate); updated Delegation Chain term to include in-place (cross-time) narrowing. §12 Future Work: additional payment-related action types, spending_limit detailed specification, and pre-commitment negotiation flexibility added as roadmap items. Companion to Protocol Specification v0.1.5; AID-field and delegation-chain detail edits in the Spec do not apply here per the Framework’s lifecycle-focused scope.
v0.1.4	May 2026	Sections 5.2 and 5.3 removed — these described AID field definitions and did not fit the framework’s basic / lifecycle-focused scope. The role/verification_tier matrix and role/authorized_actions matrix remain normative in the Protocol Specification (§5.2.1 and §5.3 Table 5.5). Revision history relocated from title page to §14.
v0.1.3	May 2026	New Section 5.3 “Authorized Actions” published — fixed enumeration (purchase, sell, hire, contract, delegate) with role compatibility matrix. Operational scenarios §10.1–10.3 illustrate example values per role.
v0.1.2	May 2026	New Section 5.2 “Principal Verification by Role” published — role/verification_tier matrix with normative rejection semantics. Operational scenarios §10.1–10.3 specify verification tier required per role.
v0.1.1	April 2026	Initial Framework Edition release.

End of AEA/P Protocol Framework v0.2.2

https://aeap.dev