Published by AI ROF™ · Version 1.0 · 28 July 2026

The Agent Estate Standard

A free, vendor-neutral management standard for AI agents in production. It defines the Agent Record, the accountabilities behind every agent, the operating cadence that keeps an estate under control, and three levels of conformance.

Vijayan Seenisamy · Author of The Pilot Trap · Enterprise Agentic AI Delivery
Licensed under Creative Commons Attribution-NoDerivatives 4.0 (CC BY-ND 4.0)
The canonical text lives at aideliverydiscipline.com/standard/v1/

Published by AI ROF™ · Version 1.0 · 28 July 2026

The Agent Estate Standard

The canonical text of version 1.0. This page is permanent: its address and its content do not change after publication. Cite requirements by their anchors.

Download PDF · View JSON schema · Licensed CC BY-ND 4.0

The Agent Estate Standard

Version 1.0 Published by AI ROF™ Author: Vijayan Seenisamy

Status of this document

This is version 1.0 of the Agent Estate Standard, published as a final specification.

The Standard is free to read, implement and cite. It will remain free.

Version 1.0 supersedes no prior publication. Subsequent revisions follow the change control rules in the final section of this document.

Licensed under Creative Commons Attribution-NoDerivatives 4.0 International (CC BY-ND 4.0). Free to read, share and cite, with attribution. No derivative versions.

Scope

This Standard defines how an enterprise manages AI agents operating in production.

In scope: the register of agents, the accountabilities attached to each one, the operating cadence that governs promotion, review and retirement, and the economics that determine whether an agent should continue to run.

Out of scope: how agents are built, which models or platforms are used, whether an individual agent is technically safe to deploy, and how agents are secured. Those questions are addressed by vendor tooling, by engineering practice, and by the risk frameworks listed in Appendix C. This Standard assumes they are answered elsewhere and addresses what happens next.

This Standard does not replace an agent registry. It defines what a registry is for.

Normative language

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

MUST indicates an absolute requirement for the stated conformance level. SHOULD indicates a recommended practice which may be departed from where the reason is recorded. MAY indicates an option.

Section 1: Definitions

1.1 **Agent** A software system that takes actions toward a goal with a degree of autonomy, where the sequence of actions is determined at runtime rather than fully specified in advance.

1.2 **Agent estate** The complete set of agents operating in production across an organisation, irrespective of platform, business unit, or who built them.

1.3 **Production** The state in which an agent acts on live data or live systems, or produces output that reaches a person or process outside the delivery team. Pilots operating on live data are in production for the purposes of this Standard.

1.4 **Promotion** The decision to move an agent into production. Promotion is a decision with evidence behind it, not an event that occurs when deployment completes.

1.5 **Retirement** The permanent withdrawal of an agent from production. Distinct from suspension, which is temporary, and demotion, which returns an agent to a supervised or assistive autonomy level.

1.6 **Successful outcome** A completed unit of work that met its acceptance condition without human rework. Where an outcome required human correction before it could be used, it is not a successful outcome for the purpose of cost measurement.

1.7 **Business Owner** The named individual accountable for an agent producing its stated outcome. Distinct from the person accountable for the agent functioning.

1.8 **Agent Record** The record held for each agent in production, defined in Section 2.

1.9 **Estate review** The recurring forum with authority to continue, remediate, demote or retire an agent, defined in Section 5.

1.10 **Conformance** The state of satisfying all requirements of a stated level, evidenced on request without preparation.

Section 2: The Agent Record

The Agent Record is the unit of management in this Standard. One record exists per agent in production.

The Record is not a replacement for a vendor agent registry. It is the management layer above it. Identity fields can generally be sourced from a registry. Accountability and governance fields cannot, because no registry currently holds them.

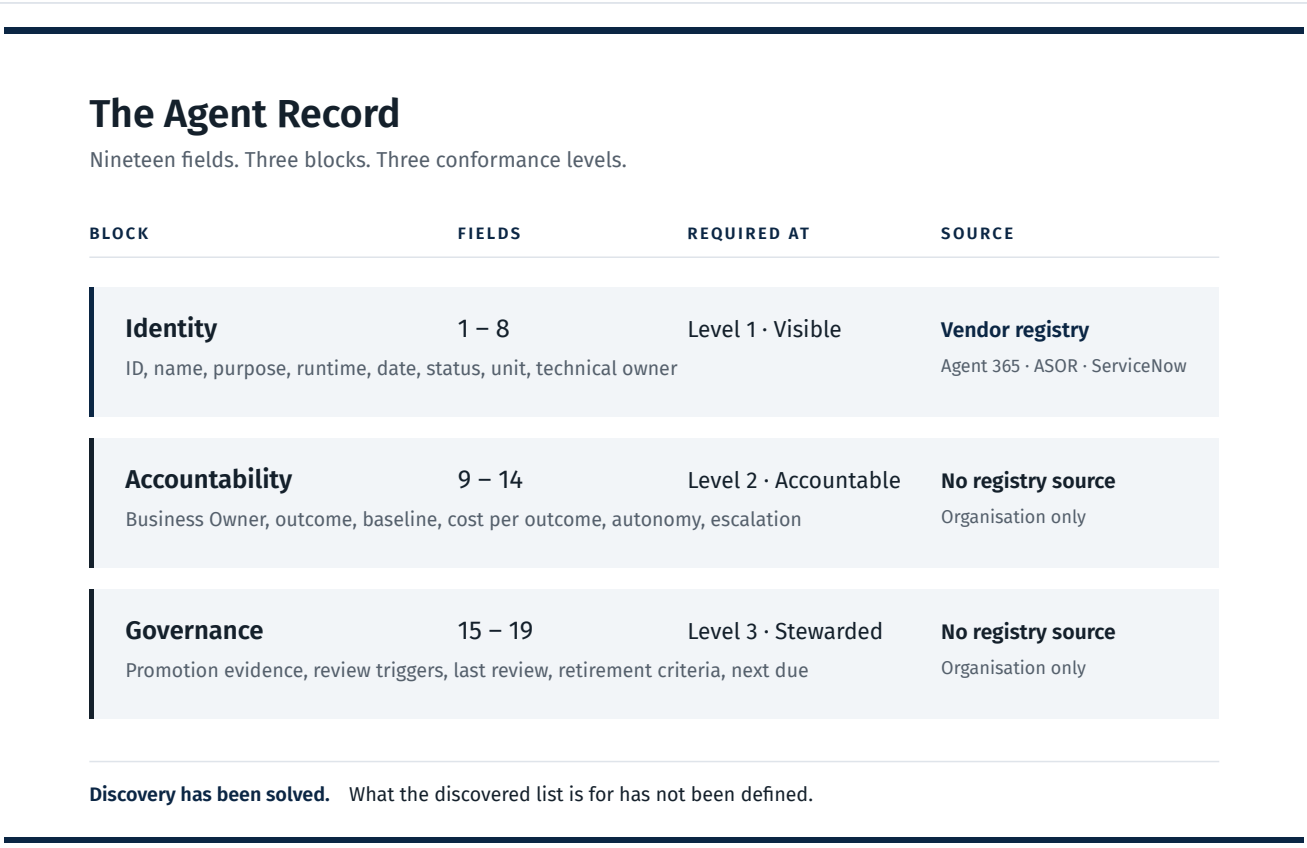


Figure 1. The Agent Record: three blocks mapped to three conformance levels.

2.1 Structure

Block	Fields	Required at
Identity	1 to 8	Level 1
Accountability	9 to 14	Level 2
Governance	15 to 19	Level 3

2.2 Identity block

1. Agent ID. A unique, persistent identifier assigned by the organisation. MUST NOT be a vendor-assigned identifier alone, because agents migrate between platforms and vendor identifiers do not survive migration.

2. Name. Human-readable, unique within the estate.

3. Purpose. One sentence stating what the agent does, written so a person outside the delivery team understands it.

4. Runtime. The platform the agent executes on. Where an agent spans runtimes, all MUST be listed.

5. Date entered production. The date the agent first met the definition in 1.3. Not the build date and not the pilot start date.

6. Status. One of: in production, under review, suspended, retired.

7. Business unit. The unit whose work the agent performs. Where an agent serves several, the accountable unit MUST be identified.

8. Technical owner. Named individual accountable for the agent functioning.

2.3 Accountability block

9. Business Owner. Named individual per 1.7. A team, role, committee or distribution list MUST NOT be recorded in this field.

10. Stated outcome. The business result the agent was funded to produce, in the terms its Business Owner would use with their own leadership.

11. Baseline. The pre-agent cost, time, volume or quality of the same outcome. Where no baseline was captured before deployment, this field MUST record "not captured" rather than a retrospective estimate.

12. Cost per successful outcome. Total cost of operation divided by count of successful outcomes over a stated period, with the measurement method recorded alongside. Operating cost includes inference, orchestration, human review, exception handling and monitoring. At Level 2 an estimate is permitted where flagged as such and produced by a consistently applied method. At Level 3 the figure MUST be measured.

13. Autonomy level. One of: assistive (proposes, human acts), supervised (acts, human approves), autonomous (acts, human notified).

14. Escalation point. Named individual or on-call function receiving escalation when the agent fails or behaves outside expectation, reachable within a stated response time.

2.4 Governance block

15. Promotion evidence reference. Pointer to the record demonstrating the agent met the criteria in 5.2.

16. Review triggers. Thresholds against the five control surfaces which, when breached, move the agent to under review without waiting for the scheduled cycle.

17. Last review date and decision. Date, and one of: continue, remediate, demote, retire.

18. Retirement criteria. The conditions under which this agent should be withdrawn. MUST be defined at promotion, not at first review.

19. Next review due. Date, set per the cadence in 5.3.

2.5 Optional fields

An organisation MAY record vendor, underlying models, cost centre, dependent agents, regulatory classification, data classification, incident history and version.

None are required. A Record carrying every optional field and lacking field 9 conforms at no level.

2.6 Sourcing the Record

Fields 1 to 8 can be substantially populated from an existing agent registry. Appendix D provides field mappings for the major platforms.

Fields 9 to 19 cannot be sourced from any registry currently available, because they record business accountability, economics and governance decisions rather than technical identity.

This division is why this Standard exists. Discovery has been solved. What the discovered list is for has not been defined.

2.7 Record integrity

2.7.1 A Record MUST be updated within ten working days of a change to any required field.

2.7.2 A Record MUST NOT be deleted on retirement. Status changes to retired and the Record is retained.

2.7.3 Where a required field cannot be populated, the Record MUST state the reason rather than remain blank. A blank field and a knowingly unavailable field are different conditions and an assessor MUST be able to distinguish them.

Section 3: The Four Roles

The Standard assigns accountability through four roles. These describe accountabilities, not necessarily four headcounts. In a small estate the same person may hold several. What cannot be compressed is the accountability itself.

3.1 **AI Product Manager.** Accountable for the business case remaining true after deployment. Owns the stated outcome, the baseline, and whether value is still being realised. Primary on Business Alignment.

3.2 **AI Architect.** Accountable for structural integrity. Data pipelines, API reliability, environment parity, rollback capability. Primary on Integration Architecture.

3.3 **AgentOps Lead.** Accountable for technical monitoring. Drift detection, reliability tracking, cost tracking per agent, override rates. Primary on Observability and Evaluation.

3.4 **AI Delivery Manager.** The system steward. The only role with sight of all five control surfaces. Primary on Governance and Compliance and on Orchestration and Operations.

3.5 **Engineering capacity** supports all four. Engineers build the integrations, evaluation harnesses and observability infrastructure. They do not own governance, business alignment or coaching decisions. That separation is what distinguishes a stewardship model from a delivery model that happens to include AI.

3.6 An organisation that assigns all five control surfaces to a single person has not adopted this Standard. It has renamed the problem.

3.7 Every agent in production MUST have each of the four accountabilities assigned to a named individual. The same individual MAY hold more than one.

Section 4: The Five Control Surfaces

Five surfaces, each giving a different form of authority over the system, all five required simultaneously. Neglect one and the system fails in ways the other four cannot detect.

Surface	What it governs
Business Alignment	Business case, cost per outcome, value realisation, sunset criteria, portfolio economics
Integration Architecture	Data pipelines, API reliability, environment parity, rollback capability, deployment automation
Governance and Compliance	Risk classification, compliance, review cadence, escalation, human override, kill switches
Observability and Evaluation	Drift detection, reliability tracking, cost per agent, adoption monitoring, override rates
Orchestration and Operations	Named ownership, delivery rhythm, cross-team coordination, agent coaching, escalation paths

4.1 Review triggers MUST be defined against each of the five surfaces for every agent in production at Level 3.

4.2 A trigger MUST be expressed as a threshold with a value, not as a qualitative condition. "Accuracy degrades" is not a trigger. "Successful outcome rate falls below 92% over a rolling 14 days" is.

4.3 An organisation MUST define at minimum one trigger per surface per agent. It MAY define more.

4.4 Triggers SHOULD be monitored automatically where the surface permits it. Business Alignment and Governance and Compliance frequently require human assessment, and that is acceptable provided the assessment cadence is defined.

Section 5: The Operating Cadence

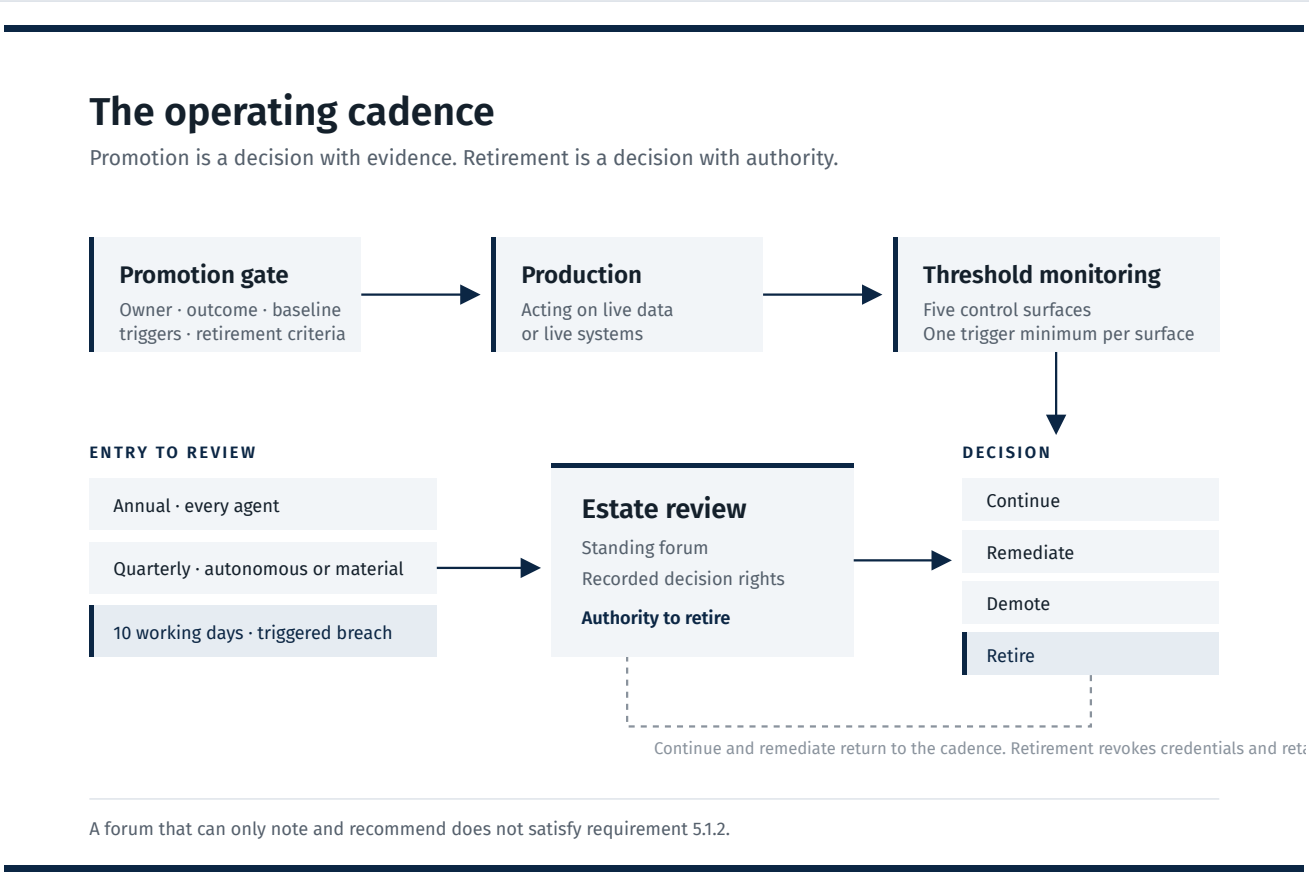


Figure 2. The operating cadence.

5.1 The estate review

5.1.1 An estate review MUST exist as a standing forum with recorded decision rights.

5.1.2 The review MUST have authority to continue, remediate, demote or retire an agent. A forum that can only note and recommend does not satisfy this requirement.

5.1.3 Attendance MUST include the AI Delivery Manager and the Business Owner of any agent under discussion.

5.1.4 Decisions MUST be recorded against the agent's Record within five working days.

5.2 Promotion gates

5.2.1 No agent MAY enter production without a promotion record.

5.2.2 A promotion record MUST evidence: a stated outcome and baseline, a named Business Owner, an assigned set of the four accountabilities, defined review triggers across all five surfaces, and defined retirement criteria.

5.2.3 Promotion MUST be a decision made by a named person, recorded with a date.

5.2.4 An agent already in production at the time an organisation adopts this Standard MUST be brought to promotion-record standard within twelve months, or retired.

5.3 Review cadence

5.3.1 Every agent in production MUST be reviewed at least annually.

5.3.2 Agents at autonomous autonomy level, or exceeding a materiality threshold the organisation defines and states, MUST be reviewed at least quarterly.

5.3.3 An agent breaching a review trigger MUST be reviewed within ten working days of the breach being detected, irrespective of its scheduled date.

5.3.4 The materiality threshold MUST be stated in terms an assessor can verify, such as annual operating cost, transaction volume, or regulatory classification.

5.4 Reconciliation

5.4.1 The register MUST be reconciled against source systems at least monthly where automated discovery is available.

5.4.2 Where automated discovery is not available for a runtime, reconciliation MUST occur at least quarterly by attested manual survey, and the runtime MUST be recorded as manually reconciled.

5.4.3 Agents discovered but not registered MUST be recorded with a discovery date and a resolution owner, and resolved within twenty working days.

5.5 Retirement

5.5.1 Retirement criteria MUST be defined at promotion.

5.5.2 An agent meeting its retirement criteria MUST be reviewed. The review MAY decide to continue, provided the reason is recorded and the criteria are revised.

5.5.3 On retirement, the agent's access and credentials MUST be revoked and the revocation recorded.

5.5.4 The Record is retained per 2.7.2.

Section 6: Estate Economics

6.1 Cost per successful outcome

6.1.1 Cost per successful outcome is total cost of operation divided by count of successful outcomes over a stated period.

6.1.2 Total cost of operation MUST include inference, orchestration, human review, exception handling and monitoring. It MUST NOT be limited to model or platform charges.

6.1.3 Outcomes requiring human rework before use MUST NOT be counted as successful.

6.2 The negative unit economics test

6.2.1 An agent is in negative unit economics when its cost per successful outcome exceeds the baseline cost of the same outcome by the process it replaced.

6.2.2 An agent in negative unit economics MUST be flagged for review at the next estate review.

6.2.3 Negative unit economics is not automatically a reason to retire. An agent MAY be retained where the reason is recorded, for example where volume growth is expected to change the position, or where the outcome carries value the cost comparison does not capture.

6.2.4 An agent retained in negative unit economics for more than two consecutive review cycles MUST have a remediation plan or a retirement date.

6.3 Attribution

6.3.1 Every unit of AI operating spend SHOULD be attributable to an agent, an owner and an outcome.

6.3.2 The proportion of AI operating spend that cannot be attributed MUST be measurable and stated.

6.3.3 Unattributed spend exceeding a threshold the organisation defines MUST be reported to the estate review.

6.4 Operating cost planning

6.4.1 Operating cost SHOULD be planned at no less than parity with build cost for the first twelve months of production, and at higher ratios for agents in regulated industries or high-volume environments.

6.4.2 An agent promoted without a funded operating budget MUST be recorded as such on its promotion record.

Section 7: Conformance

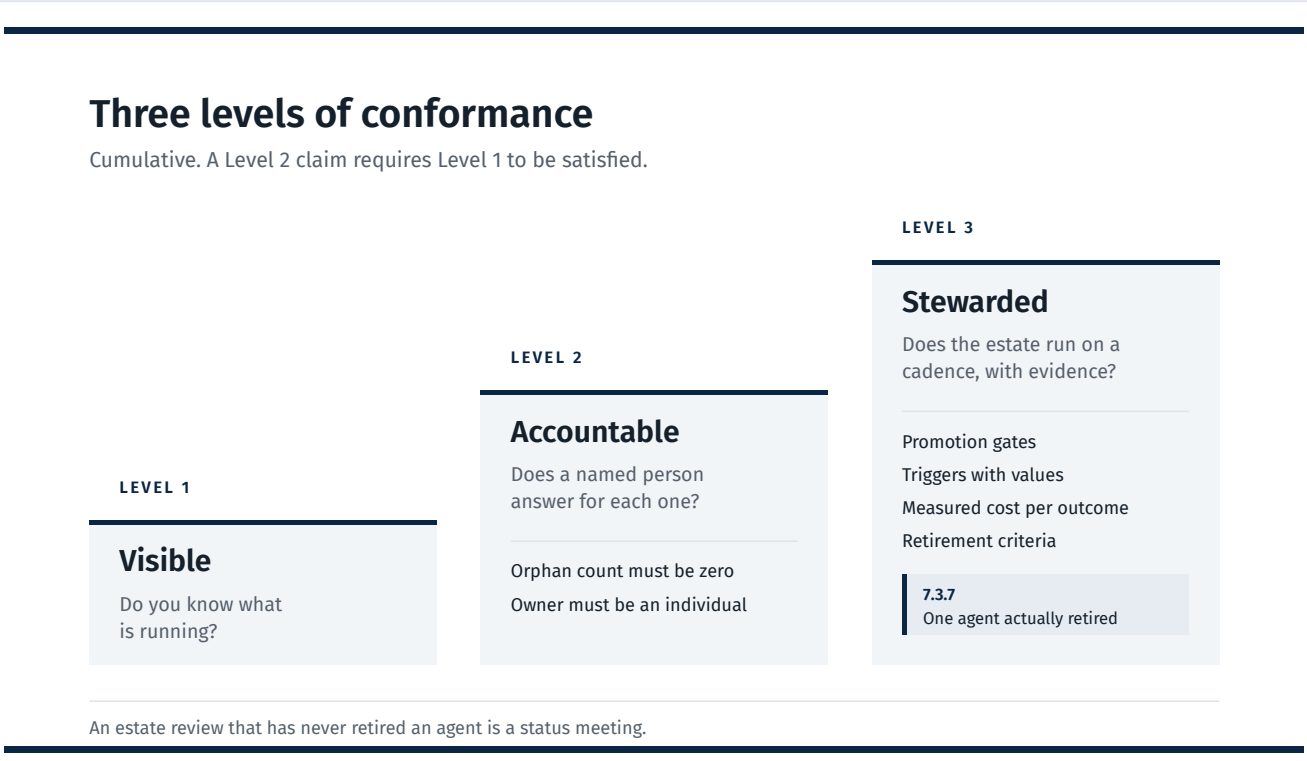


Figure 3. Three levels of conformance.

Conformance is a state held, not a report produced. An organisation conforms at a level when it can evidence that level on request without preparation.

The levels are cumulative. A Level 2 claim requires Level 1 to be satisfied.

7.1 Level 1 · Visible

The question: do you know what is running?

7.1.1 A single register MUST exist covering every agent in production, irrespective of platform.

7.1.2 Each registered agent MUST hold the identity fields in 2.2.

7.1.3 Reconciliation MUST meet the requirements of 5.4.

7.1.4 Coverage MUST be measurable and stated as a figure: agents registered divided by agents discovered.

7.1.5 Agents discovered but not registered MUST be recorded and resolved per 5.4.3.

Evidence: the register, the most recent reconciliation record with method and discovered-but-unregistered count, and a stated coverage figure.

Failure signature: the organisation produces a list but cannot say whether it is complete. A list of unknown completeness is not a register.

7.2 Level 2 · Accountable

The question: does a named person answer for each one?

7.2.1 Every agent in production MUST have a Business Owner recorded as a named individual.

7.2.2 Every agent MUST have a stated outcome with a baseline, or "not captured" recorded per field 11.

7.2.3 Every agent MUST have a cost per successful outcome produced by a stated and consistently applied method. Estimates are permitted at this level where flagged.

7.2.4 The four accountabilities MUST be assigned per 3.7.

7.2.5 The count of agents in production without a Business Owner MUST be zero.

7.2.6 Agents in negative unit economics MUST be identified and flagged per 6.2.2.

Evidence: register with owner fields at 100% coverage for production agents, outcome statements with baselines, cost figures with stated method, accountability assignments.

Failure signature: ownership recorded as a team or a platform group. Cost recorded as total programme spend rather than per outcome. Outcomes stated as efficiency or productivity with no baseline. Each produces a register that looks populated and answers nothing.

7.3 Level 3 · Stewarded

The question: does the estate operate on a cadence, with evidence?

7.3.1 No agent MAY enter production without a promotion record per 5.2.

7.3.2 The estate review MUST operate per 5.1 with recorded decision rights.

7.3.3 Review triggers MUST be defined across all five control surfaces per Section 4 and monitored between reviews.

7.3.4 The review cadence MUST meet 5.3, including the materiality threshold being defined and stated.

7.3.5 Cost per successful outcome MUST be measured rather than estimated.

7.3.6 Retirement criteria MUST be defined for every agent in production.

7.3.7 The organisation MUST be able to evidence at least one agent retired or demoted as a result of an estate review decision within the preceding twelve months.

7.3.8 Promotion, review and retirement records MUST be retained for a defined period and traceable by an assessor without reconstruction.

Evidence: promotion records, estate review minutes showing decisions rather than attendance, threshold monitoring records, and at least one executed retirement or demotion with the decision that caused it.

Failure signature: reviews take place and nothing is ever stopped. An estate review that has never retired an agent is a status meeting. Requirement 7.3.7 exists because it is the only test in this Standard that cannot be satisfied on paper.

7.4 Assessment

7.4.1 An organisation MAY self-assess and state its level publicly.

7.4.2 A self-assessment MUST state the assessment date and the assessor.

7.4.3 An independent assessment MUST be performed by a party with no delivery responsibility for the agents assessed.

7.4.4 A conformance claim MUST state the level, the assessment date, the scope assessed, and whether the assessment was self-performed or independent.

7.4.5 A claim older than twelve months SHOULD NOT be relied upon.

7.5 Scope of a claim

7.5.1 An organisation MAY claim conformance for a defined subset of its estate provided the subset is stated.

7.5.2 Permitted scoping dimensions: legal entity, business unit, geography, runtime.

7.5.3 A claim that does not state its scope MUST be read as covering the entire organisation.

Appendix A: The Five-Agent Test

Informative. A one-week entry diagnostic for organisations with no register.

Do not begin with the whole estate. Begin with five agents.

1. Select the five agents in production that matter most by cost, volume or exposure.
2. For each, hold one conversation per control surface with the person closest to it.
3. Attempt to complete fields 1 to 14 of the Agent Record.
4. Record which fields could not be completed, and why.
5. Score each agent against the five surfaces on a single page.

The output is not a maturity score. It is a list of the questions the organisation cannot currently answer about the five agents it cares most about. That list is the starting position.

Organisations completing this test typically discover they sit below Level 1.

Appendix B: Agent Record Template

Status: informative. This template implements Section 2. Where the two disagree, Section 2 governs.

Available in machine-readable form as `agent-record.schema.json`.

B.1 One-page record

For use per agent. Fields marked L1, L2 and L3 indicate the conformance level at which each becomes required.

IDENTITY

1. Agent ID `L1`	
2. Name `L1`	
3. Purpose `L1`	
4. Runtime `L1`	
5. Entered production `L1`	
6. Status `L1`	☐ In production ☐ Under review ☐ Suspended ☐ Retired
7. Business unit `L1`	
8. Technical owner `L1`	

ACCOUNTABILITY

9. Business Owner `L2`	*Named individual. Not a team, role, committee or list.*
10. Stated outcome `L2`	
11. Baseline `L2`	☐ Recorded: _____ ☐ Not captured
12. Cost per successful outcome `L2`	Figure: _____ Period: _____ Method: _____ ☐ Measured ☐ Estimated *(estimates permitted at L2 only)*
13. Autonomy level `L2`	☐ Assistive ☐ Supervised ☐ Autonomous
14. Escalation point `L2`	Name: _____ Response time: _____

Accountability assignment `L2`

Role	Named individual
AI Product Manager	
AI Architect	
AgentOps Lead	
AI Delivery Manager	

GOVERNANCE

15. Promotion evidence `L3`	**Reference: _____ Approved by: _____ Date: _____**
16. Review triggers `L3`	*One minimum per control surface. Each expressed as a value.*
17. Last review `L3`	Date: _____ Decision: ☐ Continue ☐ Remediate ☐ Demote ☐ Retire
18. Retirement criteria `L3`	
19. Next review due `L3`	

Review triggers detail `L3`

Control surface	Trigger	Threshold value	Monitoring
Business Alignment			☐ Auto ☐ Manual
Integration Architecture			☐ Auto ☐ Manual
Governance and Compliance			☐ Auto ☐ Manual
Observability and Evaluation			☐ Auto ☐ Manual
Orchestration and Operations			☐ Auto ☐ Manual

OPTIONAL
Vendor · Underlying models · Cost centre · Dependent agents · Regulatory classification · Data classification · Incident history · Version

---

B.2 Register row format

For a spreadsheet or database implementation. Column order follows field order.

```
agent_id | name | purpose | runtime | entered_production | status | business_unit | technical_owner | business_owner | stated_outcome | baseline | baseline_captured | cost_per_successful_outcome | cost_period | cost_method | cost_basis | autonomy_level | escalation_name | escalation_response_time | ai_product_manager | ai_architect | agentops_lead | ai_delivery_manager | promotion_evidence_ref | promotion_approved_by | promotion_date | review_triggers | last_review_date | last_review_decision | retirement_criteria | next_review_due | field_exceptions
```

B.3 Controlled values

Field	Permitted values
`status`	in_production, under_review, suspended, retired
`autonomy_level`	assistive, supervised, autonomous
`cost_basis`	measured, estimated
`baseline_captured`	true, false
`last_review_decision`	continue, remediate, demote, retire
`monitoring`	automatic, manual

B.4 Validation rules

These implement Section 2 and Section 7. An implementation SHOULD enforce them.

V1. Where `status = in_production`, field 9 `business_owner` MUST NOT be empty. This is requirement 7.2.5 expressed as a constraint. It is the single most important validation in the template.

V2. Field 9 MUST NOT contain a value matching a team, role, committee or distribution list pattern. Implementations SHOULD validate against the organisation's directory to confirm the value resolves to an individual.

V3. Where `baseline_captured = false`, field 11 MUST read "not captured". A retrospective estimate MUST NOT be entered.

V4. Where a conformance claim is at Level 3, `cost_basis` MUST be `measured` for every agent in production.

V5. Field 5 `entered_production` MUST NOT be inherited unverified from a vendor registry creation date. See Appendix D.

V6. Where `status = retired`, the record is retained. Deletion MUST be blocked.

V7. Any required field left empty MUST have a corresponding entry in `field_exceptions` stating the reason. Empty with no exception recorded fails 2.7.3.

V8. `next_review_due` MUST NOT exceed twelve months from `last_review_date`, or three months where `autonomy_level = autonomous` or the agent exceeds the stated materiality threshold.

B.5 Implementation note

This template can be implemented in a spreadsheet, a configuration management database, a governance platform, or as an extension to an existing agent registry.

The format does not matter. What matters is that the fields exist, the validation rules hold, and the record can be produced on request without preparation.

An organisation implementing this in a spreadsheet and enforcing V1 conforms at Level 2. An organisation implementing it in an enterprise platform without enforcing V1 does not.

Appendix C: Regulatory and Framework Mappings

Status: informative. Nothing in this appendix creates a requirement. Conformance with the Agent Estate Standard is determined solely by Section 7.

Verified: July 2026. Several frameworks referenced here revise frequently. AIUC-1 updates quarterly. The IMDA framework moved from v1.0 to v1.5 within four months of release. Verify against the current published version of any framework before relying on a mapping.

C.1 Why this appendix exists

Between January and February 2026, four significant agentic AI frameworks published within six weeks of each other. A reasonable question follows: does an enterprise that has adopted one of them need this Standard as well?

The answer turns on what each framework is for.

The frameworks listed here assess **whether an agent should be allowed to operate**. They address risk, safety, security, identity and compliance. They are the right instruments for those questions and this Standard does not attempt to replace any of them.

This Standard addresses **whether an agent is still worth operating**, and who answers for it. Accountability by name, economics per outcome, promotion evidence, review cadence, retirement.

Both questions have to be answered. Neither answers the other.

C.2 How to read the mappings

Each table shows where a requirement of this Standard supports an obligation or control area in the named framework. A mapping means the Standard's requirement produces evidence that is useful to the other framework. It does not mean satisfying one satisfies the other.

Mappings are stated at the level the source framework publishes at. Where a framework publishes numbered articles or clauses, those are cited. Where a framework publishes dimensions or pillars, those are cited.

C.3 NIST AI Risk Management Framework (AI 100-1)

Four functions: Govern, Map, Measure, Manage.

AI RMF function	Supported by
Govern	Section 3 (four roles), 5.1 (estate review with decision rights), 7.2.1 (named Business Owner), 7.2.4 (accountabilities assigned)
Map	Section 2 (Agent Record), 2.2 (identity block), 7.1 (register coverage), Section 4 (control surfaces)
Measure	4.2 (numeric thresholds), Section 6 (cost per successful outcome), 6.2 (negative unit economics test)
Manage	5.2 (promotion gates), 5.3 (review cadence), 5.5 (retirement), 7.3.7 (evidenced retirement or demotion)

What the Standard adds: AI RMF is risk-management guidance without a certification path and without a defined unit of management. This Standard supplies the unit, being the Agent Record, and the cadence that operates on it.

C.4 NIST AI Agent Standards Initiative (CAISI, February 2026)

Three pillars: industry-led standards, open-source protocol development, and security and identity research. The associated NCCoE concept paper addresses agent identity and authorisation, covering authentication, authorisation, auditing, non-repudiation and prompt injection mitigation.

Initiative area	Supported by
Agent identity	2.2 field 1 (organisation-assigned persistent ID, independent of vendor identifier)
Authorisation	2.3 field 13 (autonomy level), optional field for systems and data accessed
Auditing and non-repudiation	7.3.8 (records traceable without reconstruction), 5.1.4 (decisions recorded within five working days)
Lifecycle	5.2 (promotion), 5.5 (retirement including credential revocation)

What the Standard adds: the Initiative addresses identity and security. It does not address business accountability or economics. Field 1 of the Agent Record is deliberately specified to survive platform migration, which vendor-assigned identity does not.

C.5 ISO/IEC 42001:2023, AI Management Systems

Certifiable through accredited bodies. Targets the management system rather than agent behaviour.

ISO/IEC 42001 area	Supported by
Clause 5, leadership and roles	Section 3 (four roles), 3.7 (named individuals per agent)
Clause 6, planning and objectives	2.3 field 10 (stated outcome), field 11 (baseline), 6.4 (operating cost planning)
Clause 8, operation	Section 5 in full (promotion, review, reconciliation, retirement)
Clause 9, performance evaluation	Section 4 (control surfaces and triggers), Section 6 (estate economics), 5.1 (estate review)
Clause 10, improvement	5.5.2 (retirement criteria revision with recorded reason), CC change control
AI system lifecycle controls	5.2 (promotion gates), 5.5 (retirement)
AI system impact assessment	2.3 fields 10 to 14, 5.2.2 (promotion evidence)

What the Standard adds: ISO/IEC 42001 requires a management system and leaves the operational unit undefined. An organisation can hold 42001 certification and still be unable to name the owner of a given agent. This Standard makes the agent the unit and the owner a required field.

C.6 EU AI Act

Article numbers refer to Regulation (EU) 2024/1689. Applicability depends on risk classification and on whether the organisation is a provider or a deployer.

EU AI Act obligation	Supported by
Art. 9, risk management system	Section 4 (control surfaces), 4.1 to 4.3 (triggers per surface)
Art. 12, record-keeping and logging	Section 2 (Agent Record), 2.7 (record integrity), 7.3.8 (traceable records)
Art. 14, human oversight	2.3 field 13 (autonomy level), field 14 (escalation point), 5.1.2 (authority to demote)
Art. 15, accuracy and robustness	4.2 (numeric thresholds), 1.6 (successful outcome definition excluding reworked output)
Art. 17, quality management system	Section 5 (operating cadence), Section 3 (roles)
Art. 26, deployer obligations	7.2.1 (named Business Owner), 2.3 field 14 (escalation), 5.3 (review cadence)
Art. 72, post-market monitoring	5.3 (review cadence), 5.4 (reconciliation), 4.4 (trigger monitoring)

What the Standard adds: the Act imposes obligations without specifying the operating rhythm that discharges them. Requirement 5.3.3, review within ten working days of a detected trigger breach, is the kind of operational specificity the Act assumes and does not provide.

C.7 Singapore IMDA Model AI Governance Framework for Agentic AI

Launched 22 January 2026, revised to v1.5 on 20 May 2026. Voluntary and non-binding. Four dimensions.

IMDA dimension	Supported by
1. Assessing and bounding risks upfront	5.2 (promotion gates), 5.2.2 (promotion evidence requirements), 2.3 field 13 (autonomy level)
2. Making humans meaningfully accountable	7.2.1 (Business Owner as named individual), 3.7 (accountabilities assigned), 7.2.5 (zero orphans)
3. Implementing technical controls and processes	Section 4 (five control surfaces), 4.2 (numeric thresholds), 5.5.3 (credential revocation on retirement)
4. Enabling end-user responsibility	2.3 field 14 (escalation point with stated response time), 2.2 field 3 (purpose stated in non-specialist terms)

Alignment note. Dimension 2 is the closest external alignment to this Standard's Level 2. IMDA states the principle that humans must be meaningfully accountable. This Standard states the requirement that makes it verifiable: field 9 must hold a named individual, and a team, role, committee or distribution list must not be recorded in it.

v1.5 adds treatment of agent sprawl and multi-agent systemic risk. Sections 5.4 (reconciliation) and 7.1.4 (stated coverage figure) address the measurement side of sprawl directly.

C.8 AIUC-1

A certification standard for AI agents, issued by the Artificial Intelligence Underwriting Company. Six control domains, updated quarterly, certificates valid twelve months with quarterly technical retesting.

AIUC-1 domain	Supported by
Accountability	7.2 in full (Level 2), Section 3 (roles)
Reliability	4.2 (thresholds), 1.6 (successful outcome definition), 5.3.3 (triggered review)
Safety	5.1.2 (authority to demote or retire), 2.3 field 14 (escalation)
Security	Not addressed. Out of scope per this Standard's Scope statement.
Data and Privacy	Not addressed. Optional field for data classification only.
Society	Not addressed.

Relationship. AIUC-1 certifies the agent or the vendor supplying it. This Standard addresses the estate operating it. An enterprise could hold AIUC-1 certification for every purchased agent and still have no register, no named owners and no retirement decisions. The two operate at different levels and the overlap is limited to the Accountability and Reliability domains.

C.9 APRA CPS 230, Operational Risk Management

Applies to APRA-regulated entities in Australia. Effective 1 July 2025.

CPS 230 area	Supported by
Identification of critical operations	5.3.2 and 5.3.4 (materiality threshold defined and stated), 2.2 field 7 (business unit)
Operational risk profile	Section 4 (control surfaces), 4.1 (triggers per surface per agent)
Tolerance levels	4.2 (thresholds expressed as values), 6.2 (negative unit economics test)
Monitoring and escalation	4.4 (trigger monitoring), 2.3 field 14 (escalation point and response time)
Register of material arrangements	Section 2 (Agent Record), Section 7.1 (register with stated coverage)
Board and management accountability	Section 3 (roles), 5.1 (estate review with decision rights)

Note for regulated entities. Where agents perform work inside a critical operation, the materiality threshold required by 5.3.4 should be set to align with the entity's existing critical operations definition rather than defined independently. Two thresholds that disagree will produce two answers to the same question.

C.10 Reverse map: what remains after compliance

The practical version of this appendix. If an organisation already holds or follows the framework in the left column, the right column states what this Standard adds that the framework does not require.

If you already have	This Standard still adds
An agent registry (Agent 365, Workday ASOR, ServiceNow)	Fields 9 to 19. Business ownership by name, stated outcome and baseline, cost per successful outcome, promotion evidence, review triggers, retirement criteria. No registry holds any of these.
ISO/IEC 42001 certification	The agent as the unit of management, a named owner per agent, and cost per successful outcome. A certified management system does not require either.
NIST AI RMF adoption	A defined register, a promotion gate, a review cadence with authority to retire, and an evidenced retirement.
EU AI Act compliance	The operating rhythm. The Act states what must be true. It does not state the cadence, the forum, or who holds the decision.
IMDA MGF conformance	The verifiable form of Dimension 2. Accountability recorded as a named individual, with the orphan count required to be zero.
AIUC-1 certification	Everything above the individual agent. The estate, the register, the economics and the retirement decision.
CPS 230 compliance	Agent-level granularity inside the critical operation, and the economics test that CPS 230 does not contemplate.

C.11 What this Standard does not claim

It does not confer compliance with any framework listed here.

It does not substitute for a risk assessment, a conformity assessment, a security review or a data protection impact assessment.

It does not address whether an agent is safe, secure, lawful or fair to deploy. Those questions come first and are answered elsewhere.

It addresses what happens after the agent is running, which is the question none of the frameworks above were built to answer.

Appendix D: Vendor registry mappings

Informative. Field mappings from the major agent registries into the Identity block.

Record field	Microsoft Agent 365	Workday ASOR	ServiceNow
2 Name	Agent name	Agent name	Agent name
4 Runtime	Host product	Platform	Platform
5 Date entered production	Creation date*	Registration date*	Creation date*
6 Status	Availability status	Status	State
8 Technical owner	Creator	Creator	Assigned to

*Creation and registration dates approximate the production entry date and MUST be verified rather than accepted. An agent is frequently created well before it meets the definition in 1.3.

Fields 1, 3 and 7 require organisational assignment. Fields 9 to 19 have no registry source.

Appendix E: Glossary

Terms defined in Section 1, plus the following concepts referenced in this Standard and defined more fully elsewhere in the AI ROF™ body of work: System Stewardship, Agent Coaching, Negative Unit Economics, the Demo God Curse, Governance Theatre, the Production-Ready Threshold.

Change control

CC.1 This Standard uses two-part version numbering. The first part increments when a normative requirement is added, removed or materially altered. The second increments for clarification, correction or informative additions.

CC.2 A change to a normative requirement MUST be accompanied by a transition period of no less than twelve months, during which conformance claims against the prior version remain valid.

CC.3 Every published version remains permanently available at its own URL. Superseded versions are marked as superseded and are not withdrawn.

CC.4 A change log MUST accompany every release, stating what changed and why.

CC.5 Revisions are authored, not consensus-based. Reviewers are acknowledged and their input is recorded, and editorial decisions rest with the author.

The Agent Estate Standard v1.0 · Published by AI ROF™ · aideliverydiscipline.com/standard