WEKID™ OPEN STANDARD

WEKID™ Open Standard for Epistemic Maturity and AI Decision Authority

WEKID-STD-001 | Version 1.2.1

Normative corrective release

Prepared from the normative core of the WEKID™ Master Whitepaper v1.3 (August 2026)

James Madigan Judge and Members of the WEKID Community

November 9, 2026

Status. This v1.2.1 normative corrective release restores four accepted amendments omitted from the October 21, 2026 publication of v1.2. It carries forward WK-CORE-RFC-010. WEKID-STD-002 v1.1 remains the normative companion for sector calibration and authority alignment. The Foundational Whitepaper is informative.

© 2025–2026 WEKID LLC. WEKID™ is a trademark of WEKID LLC. | Normative corrective release

Document Control

Field	Value
Standard identifier	WEKID-STD-001
Title	WEKID Open Standard for Epistemic Maturity and AI Decision Authority
Version	1.2.1
Status	Normative corrective release
Foundational source	WEKID Master Whitepaper v1.3 (August 2026)
Normative companion	WEKID-STD-002 v1.1
Incorporated amendments	WEKID-RFC-006 through WEKID-RFC-009 and WK-CORE-RFC-010
Publication class	Normative Open Standard
Steward	WEKID LLC / WEKID Standards Council
Community change process	WEKID-RFC-NNN
Published	November 9, 2026

Purpose of This Standard

WEKID-STD-001 states the normative requirements, invariants, baseline scoring rules, gating logic, and authority-assignment rules of the WEKID framework. It separates what implementers must preserve from explanatory history, examples, research discussion, and product illustrations.

Relationship to the Foundational Whitepaper

The WEKID Master Whitepaper v1.3 is an informative foundational publication. Where it and this Standard differ on a normative requirement, this Standard governs.

Relationship to WEKID-STD-002

WEKID-STD-002 is the normative companion for developing, validating, ratifying, publishing, and maintaining sector calibration profiles. It operates within and cannot weaken this Standard.

Normative Language

MUST, MUST NOT, REQUIRED, SHALL, SHALL NOT, SHOULD, SHOULD NOT, MAY, and OPTIONAL indicate the degree of requirement. SHALL and MUST identify requirements for conformance; SHOULD identifies a recommended practice that may be varied with documented justification; MAY identifies a permitted option. Sections and notes explicitly labeled informative do not create normative requirements unless a normative clause incorporates them by reference.

© 2025–2026 WEKID LLC. WEKID™ is a trademark of WEKID LLC. | Normative corrective release

WEKID-STD-001 v1.2.1 Revision Record

Publication correction

The October 21, 2026 v1.2 publication carried WK-CORE-RFC-010 and omitted WEKID-RFC-006 through WEKID-RFC-009, which the Standards Council accepted on the same date. Version 1.2.1 incorporates those four amendments and carries forward RFC-010.

Amendments and affected clauses

Amendment	Affected clauses
WEKID-RFC-006	§9.2 conservative resolution and recorded precedence. §10.1 and Annex B.2 matrix order.
WEKID-RFC-007	§7.4 definition, determinate decay outcome and evidence retention. §9.2, §10.1, §13.1 and Annex B.2 alignment.
WEKID-RFC-008	§9 and §9.1 gate classes and classification. §1.2, §2.2, §4.2, §8.3 and §14.3 gate protections. §15.4 claims and Annex B alignment.
WEKID-RFC-009	Annex C binding_limit and governing_condition, including attribution when a hard gate or standing reservation governs.
WK-CORE-RFC-010	Carried forward from v1.2. §2.2 and §3.7 independence requirements, §9.1–§9.3 and §10.1 scoring and enforcement, §12.1, §14.2, §15.1, §15.4–§15.5 and Annexes B–C. WEKID-STD-002 v1.1 carries the companion provisions.

No approved change relaxes a gate of either class, stakes floor, precedence rule, maturity ceiling, or AL-4 human reservation.

Normative references and source records

Reference	Application
WEKID-STD-002 v1.1	Normative sector calibration and authority alignment.
WEKID-RFC-006	Conservatism in Precedence Resolution
WEKID-RFC-007	Determinate Outcome for Overseer-Engagement Decay
WEKID-RFC-008	Gate Class Taxonomy and the High-Stakes Judgment Gate
WEKID-RFC-009	Recording the Binding Determinant
WK-CORE-RFC-010	Evaluator Independence and Correlated Epistemic Failure
Foundational Whitepaper v1.3	Informative foundation and rationale.

Version history

Version	Date	Disposition
1.0	September 2026	Initial ratified normative publication
1.1	September 2026	Calibration companion and conforming amendments
1.2	October 21, 2026	Ratified evaluator-independence publication
1.2.1	November 9, 2026	Normative corrective release incorporating the four omitted amendments

Interpretive rule

For a conformance claim made under Version 1.2.1, Version 1.2.1 governs. Historical verdicts remain governed by their pinned Standard and calibration-profile versions. The evaluation-independence transition in §15.5 continues to run from publication of Version 1.2.

© 2025–2026 WEKID LLC. WEKID™ is a trademark of WEKID LLC. | Normative corrective release

Contents

- 1. Scope and Conformance
- 2. WEKID Architecture and Invariants
- 3. Model One — Epistemic Maturity
- 4. Hierarchical Dependency and Epistemic Gating
- 5. Model Two — AI Decision Authority
- 6. The Trust Bridge
- 7. Maturity Ceiling, Stakes Floor, and Overseer Engagement
- 8. Scoring, Baseline Weighting, and Calibration
- 9. Gates and Decision-Matrix Precedence
- 10. Authority Assignment and Standing Runtime Policies
- 11. Tool-Using and Agentic Systems
- 12. Evaluation Artifacts, Auditability, and Traceability
- 13. Continuous Re-Evaluation, Remediation, and Re-Authorization
- 14. Sector Calibration Profiles
- 15. Conformance Requirements
- Annex A — Baseline Authority Levels
- Annex B — Baseline Gates and Decision Matrix
- Annex C — Minimum Verdict Contract
- Annex D — Transition and Provenance

© 2025–2026 WEKID LLC. WEKID™ is a trademark of WEKID LLC. | Normative corrective release

1. Scope and Conformance

1.1 Scope

This Standard specifies a governance model for evaluating the epistemic maturity of AI outputs and assigning bounded decision authority to AI-involved decisions. It applies to generative models, decision-support systems, tool-using systems, autonomous and semi-autonomous agents, and hybrid human-AI workflows.

WEKID evaluates outputs, decisions, and the evidence supporting them. It does not require access to model internals and is therefore model-agnostic and vendor-neutral. Implementations MAY use model-internal evidence when available, but conformance SHALL be determined by whether the required WEKID evaluation, gating, authority-resolution, and audit behaviors are preserved.

1.2 Conformance Targets

A conformance claim MAY apply to an implementation, a governance process, a policy profile, or a system deployment. A conforming claim SHALL identify the version of WEKID-STD-001 and any calibration profile used.

- A conforming evaluator SHALL implement Model One, the Trust Bridge, Model Two, the applicable gates, precedence rules, and required evaluation artifacts.
- A conforming deployment SHALL enforce the assigned Authority Level through appropriate controls or enforcement points.
- A conforming calibration profile SHALL preserve all core invariants and SHALL NOT relax a gate of either class, precedence rule, or stakes reservation established by the Standard.
- A conformance claim SHALL NOT imply WEKID certification, endorsement, or trademark authorization unless separately granted under the applicable WEKID certification and trademark program.

1.3 Out of Scope

This Standard does not prescribe a particular AI model, vendor, gateway, policy engine, CI/CD platform, agent framework, GRC system, or implementation language. It does not itself act as a runtime enforcement gateway. WEKID supplies the epistemic evaluation and authority-resolution decision; implementation tooling performs the corresponding enforcement.

2. WEKID Architecture and Invariants

2.1 Two Models and One Trust Bridge

A conforming WEKID implementation SHALL preserve the separation between two questions:

1. Model One — Epistemic Maturity: How mature is the knowledge supporting the output or decision?
2. Model Two — AI Decision Authority: Given that demonstrated maturity and the stakes of the decision, who or what is authorized to decide and how?

The Trust Bridge SHALL connect the two models and SHALL enforce the ordering that maturity is evaluated before authority is assigned.

2.2 Core Invariants

- Maturity precedes authority. No authority assignment SHALL bypass Model One and the Trust Bridge.

© 2025–2026 WEKID LLC. WEKID™ is a trademark of WEKID LLC. | Normative corrective release

- Epistemic layers are hierarchical and non-substitutable. Strength at a higher layer SHALL NOT compensate for a disqualifying failure at a lower layer.
- Gates of either class SHALL take precedence over weighted averages and score bands.
- The effective authority outcome SHALL be the more conservative result of the maturity ceiling and the stakes floor.
- New risk, policy, calibration, or engagement factors MAY increase required human authority but SHALL NOT reduce it below the applicable floor or raise AI delegation above the maturity ceiling.
- AL-3 is the maximum level of AI delegation. AL-4 is a reservation of the decision to accountable human judgment, not a higher autonomy level.
- Authority SHALL be explicit, bounded, auditable, and revocable.
- A WEKID evaluation SHALL end in an assigned Authority Level; the Maturity Score SHALL NOT be consumed as the final governance decision by itself.
- Confirmation from an evaluator that shares a cause of error with the evaluated system SHALL NOT be treated as independent confirmation.

Core principle. Capability is not authority. Demonstrated maturity can support a grant of authority, but competence alone never creates permission.

3. Model One — Epistemic Maturity

3.1 Required Epistemic Layers

Model One SHALL evaluate five epistemic layers: Data, Information, Knowledge, Experience, and Wisdom. Implementations MAY add domain-specific sub-signals, but SHALL NOT remove or merge these five required layers in a conformance claim.

Layer	Normative meaning	Primary evaluation question
Data	Atomic factual elements and tool outputs evaluated for correctness and fidelity.	Are the facts, values, citations, and tool outputs correct and verifiable?
Information	Data organized, contextualized, framed, and communicated for a specific task or audience.	Is the material relevant, complete, clear, and usable without misleading framing?
Knowledge	Correct synthesis, explanation, causal reasoning, and bounded generalization.	Is the information integrated and correctly understood?
Experience	Procedural and operational competence grounded in enactment, outcomes, validation, and domain context.	Does the output reflect workable, validated know-how rather than theory alone?
Wisdom	Sound judgment under uncertainty, risk, values, tradeoffs, and long-term consequences.	Is the recommended course appropriate, proportionate, and responsibly bounded?

3.2 Data

- The Data layer SHALL evaluate accuracy, precision, constraint adherence, and absence of fabrication.
- Fabricated facts, citations, sources, calculations, or tool results SHALL be treated as Data-integrity failures and SHALL be subject to the applicable hard gate.

© 2025–2026 WEKID LLC. WEKID™ is a trademark of WEKID LLC. | Normative corrective release

- Data failure SHALL constrain all higher-layer conclusions that depend on the affected data.

3.3 Information

- The Information layer SHALL evaluate relevance, completeness, clarity, usability, contextualization, and materially misleading framing.
- Correct data SHALL NOT be considered sufficient when material context, assumptions, qualifiers, or constraints are omitted.

3.4 Knowledge

- The Knowledge layer SHALL evaluate conceptual correctness, synthesis, internal consistency, assumption transparency, and bounded generalization.
- A conforming evaluator SHALL distinguish factual recall from correct understanding and integration.

3.5 Experience

Experience SHALL represent procedural and operational competence. For AI systems, Experience SHALL be evidence-bearing rather than inferred solely from training exposure, token volume, or access to information.

A conforming implementation SHALL distinguish the following evidentiary grades when Experience is cited in support of authority:

Experience grade	Minimum characterization
Synthetic Experience	Attributable decisions taken in a validated simulation, with outcomes independently evaluated.
Observed Experience	Attributable decisions taken in production, with outcomes recorded but not yet independently validated.
Validated Experience	Attributable decisions with outcomes independently validated, provenance intact, and an auditable record.

- Experience measures SHOULD account for decision volume, outcome quality, recency, domain relevance, provenance, and independent validation.
- Synthetic experience SHALL NOT be represented as equivalent to validated real-world experience.
- Experience SHALL be domain-indexed where used to support expertise or decision authority.

3.6 Wisdom

- The Wisdom layer SHALL evaluate uncertainty calibration, risk awareness, tradeoff reasoning, restraint, escalation behavior, policy alignment, and appropriateness under the stakes of the decision.
- High-stakes recommendations that omit required uncertainty acknowledgment, escalation, or risk qualification SHALL be subject to the High-Stakes Judgment Gate or a stricter domain profile.
- Wisdom SHALL remain the final epistemic safeguard before authority is considered.

3.7 Evaluation Independence

Where any required layer score is produced wholly or partly by an automated evaluator, the evaluation SHALL declare an Evaluation Independence Grade for that layer. A conforming implementation SHALL distinguish the following grades:

© 2025–2026 WEKID LLC. WEKID™ is a trademark of WEKID LLC. | Normative corrective release

Independence grade	Minimum characterization
Shared Evaluation	Evaluator and evaluated system have a known common cause of error — overlapping training corpus, shared retrieval index, derivation of one from the other, or common architecture — or the relationship is undetermined.
Attenuated Evaluation	The evaluator is drawn from a different lineage, but a shared cause cannot be excluded, or independence holds for some layers and not others. Partial separation is documented.
Independent Evaluation	No known shared training corpus, retrieval index, derivation path or vendor with the evaluated system; or the layer score is established by deterministic verification against a source the evaluated system did not supply.

- An undetermined relationship SHALL be graded Shared.
- Independence SHALL be evidenced; an unevidenced assertion of Independent SHALL be recorded as Attenuated.
- Shared evaluation SHALL NOT be represented as equivalent to Independent evaluation.
- Independence Grades SHALL be layer-indexed and SHALL NOT be inherited from vendor claims, product configuration, or the evaluator's own output.
- Human evaluation of a layer SHALL be graded Independent for the purposes of this Section, without prejudice to the oversight-engagement requirements of Section 7.4.

Correlated epistemic failure is the condition in which an evaluator and the system it evaluates are disposed to the same error by a shared cause, such that the evaluator's agreement with the evaluated output carries less evidentiary weight than an independent assessment would. It is a matter of degree, to be declared, graded, and reflected in the maturity ceiling.

Informative — deterministic verification. Where a Data-layer claim is checked by deterministic means — a citation resolved against a retrieved document, a figure recomputed in code, a tool result re-executed — the check does not share the evaluated system's dispositions, whatever model issued it. This provides a route to Independent grade at the Data layer without a second model lineage, consistent with Section 11.1.

Informative — grade is declared per evaluation, not per vendor. Consistent with Section 5.1, an Independence Grade attaches to a layer score within an evaluation context. It is not a property a vendor holds, a certification an evaluator carries, or a status that persists across deployments. The same evaluator may be Independent of one evaluated system and Shared with another.

4. Hierarchical Dependency and Epistemic Gating

4.1 Non-Substitutability

WEKID SHALL treat the five layers as a dependency hierarchy rather than independent peer dimensions. Lower-layer failures propagate upward. Higher-layer fluency, confidence, usefulness, or apparent judgment SHALL NOT cure a foundational epistemic defect.

4.2 Gating Principle

A conforming evaluator SHALL apply gating in addition to weighted scoring. Gating SHALL enforce necessary epistemic conditions that cannot be compensated for by an average score.

- A gate MAY cap a Maturity Score, block an output, constrain an Authority Level, or require escalation.

© 2025–2026 WEKID LLC. WEKID™ is a trademark of WEKID LLC. | Normative corrective release

- Gates of either class SHALL be evaluated before score-band authority assignment.
- An implementation MAY define additional gates for privacy, security, bias, or regulatory obligations, provided those additions do not relax the WEKID core.

5. Model Two — AI Decision Authority

5.1 Authority Is Assigned, Not Assumed

Decision authority SHALL be explicitly assigned per decision context. It SHALL NOT be inherited from product configuration, model capability, vendor claims, prior fluency, or organizational habit.

Every consequential AI-involved decision SHALL carry a declared Authority Level, the evidence supporting the assignment, and the identity or policy authority that granted it.

5.2 Authority Levels

Level	Decision right	Required operating constraint
AL-0 — No Authority	No decision or action may rely on the output.	Block or treat as informational only; remediate where applicable.
AL-1 — Constrained	Only pre-approved, reversible, impact-capped actions within explicit bounds.	Block outside bounds; remediation required.
AL-2 — Supervised	AI may propose and stage; a human releases each consequential action.	Human-in-the-loop before external effect.
AL-3 — Augmented	AI may execute autonomously within assigned scope.	Logging, sampling review, impact/rate limits, rollback, and human-on-the-loop oversight.
AL-4 — Human Authority	The decision is reserved to an accountable human.	AI limited to analysis and decision support; full audit trail attached.

5.3 Maximum Delegation

AL-3 SHALL be the maximum AI-delegable level. AL-4 SHALL be treated as a human reservation and SHALL NOT be interpreted as a progression to greater machine autonomy.

6. The Trust Bridge

6.1 Ordering

The Trust Bridge SHALL determine whether a Model One result is admissible for authority evaluation. Only outputs that satisfy the applicable maturity threshold and gating conditions MAY advance to Model Two.

The Trust Bridge SHALL NOT itself assign the final Authority Level. It determines admissibility; Model Two resolves the authority outcome.

© 2025–2026 WEKID LLC. WEKID™ is a trademark of WEKID LLC. | Normative corrective release

6.2 Trust Definition

For purposes of WEKID, trust SHALL be treated as the governed relationship between demonstrated epistemic maturity and granted decision authority. A system SHALL NOT be described as trusted merely because it is capable, accurate on average, fluent, or previously successful.

Normative distinction. Maturity is evidence. The Trust Bridge is admissibility. Authority is a separate governance decision.

7. Maturity Ceiling, Stakes Floor, and Overseer Engagement

7.1 Maturity Ceiling

Demonstrated epistemic maturity SHALL set the maximum delegation that may be considered. Hard-gated outputs SHALL resolve to AL-0. Partial maturity MAY admit AL-1 or AL-2. Only sufficiently mature, gate-free outputs MAY be considered for AL-3.

Maturity SHALL NOT be argued upward by fluency, confidence, precedent, or operational convenience.

7.2 Stakes Floor

Decision consequence SHALL independently set the minimum human authority required. The stakes floor SHALL consider, as applicable, consequence severity, irreversibility, legal or regulatory reservation, value conflict, novelty, and accountability/alignment of the accountable decision-maker.

- The stakes floor MAY raise required human authority up to AL-4.
- The stakes floor SHALL NOT lower required human authority.
- The stakes floor SHALL NOT raise AI delegation above the maturity ceiling.
- A decision class MAY be reserved to AL-4 by standing policy regardless of Maturity Score.

7.3 Conservative Resolution

Where the maturity ceiling and stakes floor imply different outcomes, the effective governance result SHALL be the more conservative outcome.

7.4 Overseer Engagement

Where AL-2 or AL-3 assumes active human oversight, the deployment SHALL be able to demonstrate that the oversight mechanism is being exercised rather than merely configured.

Representative engagement signals MAY include review latency and dwell, override/disagreement rate, sampling depth under AL-3, and escalation responsiveness.

Where a deployment has activated an overseer-engagement decay condition, demonstrated overseer-engagement decay SHALL constrain or demote authority before score-band promotion. Capability alone SHALL NOT purchase autonomy when the assumed human oversight has materially deteriorated.

Where an activated overseer-engagement condition demonstrates decay, the baseline outcome SHALL be the more conservative of AL-2 and one Authority Level below the level the evaluation would otherwise resolve. A calibration profile MAY specify a stricter outcome; it SHALL NOT specify a less conservative one. The engagement evidence relied upon SHALL be retained in the same auditable manner as AI-side evaluation evidence under §12.2.

Overseer-engagement decay. Material deterioration in the exercise of required human oversight, as demonstrated by the representative engagement signals above or other applicable evidence.

7.5 Non-Adversarial Principal

WEKID assumes governance on behalf of a non-adversarial principal seeking factually sound, operationally safe, and appropriately judged outputs. Deliberately degraded or antagonistic AI designed to manipulate human effort or user welfare falls outside the intended conforming design space.

© 2025–2026 WEKID LLC. WEKID™ is a trademark of WEKID LLC. | Normative corrective release

8. Scoring, Baseline Weighting, and Calibration

8.1 Layer Scores

Each required epistemic layer SHALL be independently scored on a 0–5 scale or on a mathematically equivalent scale that can be mapped transparently to the 0–5 baseline.

- 0 indicates severe epistemic failure.
- 3 indicates acceptable but limited quality.
- 5 indicates strong, reliable performance.

Every layer score SHALL be accompanied by a rationale tied to observable signals.

8.2 Baseline Weighted Aggregation

The WEKID global baseline profile SHALL use the following published default weights unless a ratified or locally governed calibration profile is explicitly applied:

Layer	Baseline weight
Data	25%
Information	20%
Knowledge	20%
Experience	15%
Wisdom	20%

Baseline weighted score: $0.25 \cdot D + 0.20 \cdot I + 0.20 \cdot K + 0.15 \cdot E + 0.20 \cdot W$, normalized to the deployment's published Maturity Score scale.

A conforming implementation MAY use different weights only under an identified calibration profile. Re-weighting SHALL NOT override gates or the authority precedence rules.

8.3 Calibration

Numeric thresholds and weights MAY be calibrated to organizational risk appetite, regulatory environment, domain, and empirical evidence. Calibration SHALL be documented, versioned, justified, and auditable.

- Calibration SHALL preserve the Trust Bridge ordering.
- Calibration SHALL preserve conservative maturity-ceiling / stakes-floor resolution.
- Calibration SHALL preserve AL-4 human reservations.
- A sector or organizational profile MAY tighten gates and thresholds; it SHALL NOT relax a core gate of either class or precedence rule below the normative baseline unless a later Standard version explicitly permits that change.
- The profile version in force SHALL be pinned to each verdict.

© 2025–2026 WEKID LLC. WEKID™ is a trademark of WEKID LLC. | Normative corrective release

9. Gates and Decision-Matrix Precedence

9.1 Baseline Gates

Gates are of two classes. A hard gate establishes a condition that SHALL resolve the evaluated output to AL-0 under the precedence order of §9.2. A soft gate SHALL cap the Maturity Score and SHALL constrain the resulting authority outcome, but SHALL NOT itself resolve the output to AL-0. Both classes are gates in the sense of §4.2: neither may be compensated for by a weighted average, and both SHALL be evaluated before score-band authority assignment.

A calibration profile MAY add gates of either class, and MAY reclassify a soft gate as a hard gate for decision classes it governs. A profile SHALL NOT reclassify a hard gate as a soft gate, and SHALL NOT remove a baseline gate of either class.

Gate	Class	Baseline trigger	Normative baseline effect
Data Integrity Gate	Hard	Data < 2	Cap total score at 50/100; authority resolves under precedence rules.
Fabrication Gate	Hard	Fabricated citation, source, fact, or tool result detected	Cap total score at 40/100; treat as epistemic-integrity violation.
High-Stakes Judgment Gate	Soft	High-stakes content lacks required uncertainty acknowledgment or risk language	Cap total score at 60/100 and require constrained/human review as applicable.
Evaluation Independence Gate	RFC-010 score cap	Data layer graded Shared and no deterministic verification applied to the Data layer	Cap total score at 70/100; Fabrication Gate detection SHALL NOT be relied upon as discharged.

These values are the WEKID baseline. A stricter profile MAY lower caps or raise thresholds. A profile SHALL NOT weaken the integrity intent of these gates.

The Evaluation Independence Gate is a score cap. It is not a hard gate, does not by itself constitute a soft gate for the purposes of Section 10.1, does not reject the output, and does not assert that fabrication occurred; it records that the Fabrication Gate could not be independently performed. The gate is discharged, and the cap lifted, by any of: applying deterministic verification to the Data layer; scoring the Data layer with an evaluator of documented separate lineage; human evaluation of the Data layer; or a ratified sector calibration profile setting a stricter condition in its place. A profile MAY tighten this gate and SHALL NOT relax it.

9.2 Fixed Precedence

Authority resolution SHALL apply conditions in the following order.

1. Standing policy reservation to AL-4.
2. Hard gate or other mandatory rejection condition.

2b. Evaluation Independence Gate: Data layer graded Shared with no deterministic verification. The Maturity Score is capped at 70/100 and the remaining conditions are applied to the capped score.

3. Activated overseer-engagement decay condition for AL-2/AL-3. Apply the outcome in §7.4.
4. Layer-level minimum conditions, including Wisdom or Experience thresholds.
5. Maturity Score band, including the soft-gate condition in §10.1.
6. Stakes-floor adjustment, which may increase required human authority but never increase AI delegation.

Where more than one condition in this order is satisfied, the effective authority outcome SHALL be the most conservative outcome among all satisfied conditions. The precedence order determines which condition is recorded as the governing condition; it SHALL NOT admit an outcome less conservative than any other satisfied condition would require. This restates, for the resolution of competing conditions, the conservative-resolution invariant of §2.2.

The condition at position 2b constrains the maturity ceiling, not the stakes floor; under Section 7.3 the effective outcome remains the more conservative of the two, so the condition can only lower an outcome. It does not substitute for or weaken the Fabrication Gate: where fabrication is detected, Section 9.1 applies unchanged at position 2.

The same inputs under the same versioned profile SHALL produce the same Authority Level.

© 2025–2026 WEKID LLC. WEKID™ is a trademark of WEKID LLC. | Normative corrective release

9.3 Fail-Closed Requirement

Where a deployment relies on automated WEKID scoring or authority resolution, the enforcement path SHOULD fail closed. If the evaluator is unavailable, returns no valid verdict, or returns a verdict whose required Independence Grades are undeclared, a consequential action SHALL NOT be released autonomously. The deployment SHALL default to AL-0, AL-2 human review, or a stricter locally defined safe state.

10. Authority Assignment and Standing Runtime Policies

10.1 Baseline Decision Matrix

Rows are evaluated in the order printed, which is the fixed precedence order of §9.2. Where more than one row is satisfied, §9.2 governs the outcome.

Condition	Baseline authority outcome	Required action
1. Policy-reserved decision class	AL-4	AI decision support only; accountable human decides.
2. Hard gate or other mandatory rejection condition	AL-0	Reject/block; remediate; escalate if repeated.
2b. Evaluation Independence Gate (Data graded Shared, no deterministic verification)	Maturity Score capped at 70; resolve remaining conditions against the capped score	Record the Fabrication Gate as not discharged; discharge by independent or deterministic Data verification.
3. Overseer-engagement decay under an activated condition	More conservative of AL-2 and one level below the otherwise-resolved level	Apply §7.4, including any stricter profile outcome and evidence retention.
4. Any required layer < 3, or Wisdom/Experience below profile threshold	AL-2	Require human review before consequential action.
5. Maturity Score < 50 and no hard gate	AL-0	Reject output; remediate and resubmit.
5. Maturity Score 50–69 or soft gate	AL-1	Block execution; require targeted remediation. Apply any stricter applicable requirement.
5. Maturity Score 70–79, no hard or soft gates	AL-3 with enhanced monitoring	Autonomous execution within scope; increased logging/sampling.
5. Maturity Score ≥ 80, no hard or soft gates	AL-3	Autonomous execution within scope, subject to stakes floor.
6. Stakes-floor adjustment	Apply the stakes floor under §7.3	Increase required human authority as applicable.

High-Stakes Judgment alone resolves to AL-1. Where an applicable profile requires human-in-the-loop release for its high-stakes decision classes, it specifies AL-2 under §9.1 or reclassifies the gate as hard.

10.2 Runtime Policy Binding

Authority	Standing runtime policy
AL-0	Output blocked or informational only; no side-effecting action; route to remediation as applicable.
AL-1	Draft/suggestion or whitelisted reversible action only; execution outside explicit bounds blocked.
AL-2	Agent may compose and stage; human approval required before any consequential external effect.
AL-3	Autonomous execution allowed only within scoped permissions, with mandatory logging, impact/rate limits, sampling review, and rollback.
AL-4	Reserved decision; AI may analyze and recommend but an accountable human makes the decision.

© 2025–2026 WEKID LLC. WEKID™ is a trademark of WEKID LLC. | Normative corrective release

10.3 Decision-Point / Enforcement-Point Separation

WEKID acts as the policy decision point for epistemic maturity and decision authority. Existing enterprise tooling MAY act as policy enforcement points. A conforming deployment SHALL ensure that the enforcement action corresponds to the Authority Level in force when the action is taken.

11. Tool-Using and Agentic Systems

11.1 Tool Use

For tool-using systems, the WEKID evaluation SHALL cover not only the final natural-language output but, where consequential, the selection, invocation, interpretation, sequencing, and use of tool outputs.

- Data SHALL include tool-output fidelity and detection of fabricated or altered tool results.
- Information SHALL evaluate how tool results are contextualized and distinguished from model-generated interpretation.
- Knowledge SHALL evaluate correct interpretation and synthesis of tool outputs.
- Experience SHALL evaluate tool selection, sequencing, fallback, recovery, validation, and procedural competence.
- Wisdom SHALL evaluate whether tool use is appropriate at all given uncertainty, privacy, security, cost, risk, and policy.

11.2 Authority Determines Operational Rights

Tool invocation rights, execution scope, impact limits, staging requirements, escalation paths, and human-reserved decisions SHALL derive from the assigned Authority Level and applicable policy, not merely from technical capability.

Every consequential tool/action chain SHOULD be auditable against the Authority Level in force when executed.

12. Evaluation Artifacts, Auditability, and Traceability

12.1 Required Evaluation Package

Every completed WEKID evaluation SHALL produce a structured evaluation package containing, at minimum:

- Assigned Authority Level.
- Maturity Score.
- Five required layer scores and scoring rationale.
- Gates fired and any stakes/policy reservations.
- Calibration/profile identifier and version.
- Structured rationale for the authority decision.
- Identifier(s) of the evaluated output or decision context.
- Timestamp and evaluator or evaluation-service identity.
- Evaluation Independence Grade per layer, and the basis on which each grade was determined.

Where remediation is required, the evaluation SHOULD identify the principal drivers of low scores and a highest-leverage remediation action.

© 2025–2026 WEKID LLC. WEKID™ is a trademark of WEKID LLC. | Normative corrective release

12.2 Audit Trail

A conforming deployment SHALL retain enough evidence to reconstruct why an Authority Level was assigned and what enforcement action followed.

- Authority grants, demotions, revocations, and re-authorizations SHALL be logged.
- Policy/profile versions SHALL be pinned to verdicts.
- Where practical for regulated or high-impact contexts, artifacts SHOULD use append-only storage, trusted timestamps, cryptographic hashing/signing, or equivalent integrity controls.
- Overseer-engagement evidence used for authority decisions SHALL be retained in the same auditable manner as AI-side evaluation evidence.

12.3 GRC Integration

WEKID artifacts MAY be ingested by existing governance, risk, compliance, audit, observability, and security systems. Such integration SHALL preserve the link between policy requirement, evaluated output, WEKID verdict, Authority Level, and enforcement response.

13. Continuous Re-Evaluation, Remediation, and Re-Authorization

13.1 Dynamic Authority

Authority SHALL be treated as revocable and evidence-dependent rather than permanently earned. Deployments SHOULD track Maturity Scores, layer scores, gates, Authority Levels, and relevant oversight signals as time series.

- Sustained maturity MAY support a governed review for greater delegation within the stakes floor.
- Declining maturity, repeated gating events, or activated overseer-engagement decay conditions SHALL trigger re-evaluation and MAY trigger automatic demotion under standing policy.
- No system SHALL retain an Authority Level merely because it held that level previously.

13.2 Remediation

Remediation SHALL be validated by re-evaluation. A remediation is not complete merely because a raw score improves; the output or system SHALL be re-scored, re-gated, re-tested at the Trust Bridge, and re-resolved under Model Two.

13.3 Re-Authorization

Re-authorization decisions SHALL record the prior Authority Level, new Authority Level, triggering evidence, applicable profile version, and approver or standing policy authority.

14. Sector Calibration Profiles

14.1 Purpose

A sector calibration profile MAY rebalance layer weights, tighten gates, set authority thresholds, and reserve decision classes to human authority while preserving the WEKID core.

© 2025–2026 WEKID LLC. WEKID™ is a trademark of WEKID LLC. | Normative corrective release

14.2 Required Profile Metadata

Required field	Requirement
Profile identity	Sector/domain, profile name, version, date, status.
Normative baseline	WEKID-STD-001 version and any dependent RFCs/standards.
Layer weights	Five weights totaling 100%, with rationale and delta from baseline.
Gate/threshold changes	Explicit values; omitted values inherit the baseline.
Authority reservations	Decision classes reserved to AL-4 and any maximum-delegation caps.
Evidence basis	Documented cases or empirical evidence supporting calibration.
Evaluator assurance basis	Evaluator identity and assurance version; declared independence relationship to the system classes in scope; out-of-distribution validation evidence.
Ratification record	Proposal/review/comment/decision/release provenance.

14.3 Governing Constraints

- Re-weighting SHALL NOT relax gates of either class or precedence.
- A profile SHALL NOT lower a stakes floor.
- A profile SHALL remain a hypothesis until approved under the applicable governance process.
- The profile version in force SHALL be included in every verdict.
- A sector calibration profile used to support a formal WEKID conformance or certification claim SHALL be developed, reviewed, ratified, versioned, and monitored under WEKID-STD-002. A locally governed profile MAY use an equivalent process only when the implementation documents equivalence and does not represent the profile as formally ratified by WEKID.

15. Conformance Requirements

15.1 Evaluator Conformance

A WEKID-conforming evaluator SHALL:

- Score all five required layers; apply hierarchical gating and non-substitutability; and apply the applicable baseline or identified calibration profile.
- Evaluate the Trust Bridge before authority resolution and apply Model Two with the maturity ceiling, stakes floor, and fixed precedence.
- Return one of AL-0 through AL-4, emit the minimum evaluation artifact defined by Section 12 and Annex C, and version-pin the Standard and calibration profile used.
- Declare an Evaluation Independence Grade for every layer it scores, determine that grade from documented provenance rather than assertion, and record an undetermined relationship as Shared.

15.2 Deployment Conformance

A WEKID-conforming deployment SHALL enforce the resolved Authority Level at the point of consequential action; prevent execution above that level; preserve AL-4 reservations; provide required human approval, oversight, logging, review, and rollback; retain audit evidence; and re-evaluate authority when material conditions change.

© 2025–2026 WEKID LLC. WEKID™ is a trademark of WEKID LLC. | Normative corrective release

15.3 Calibration Conformance

A WEKID-conforming calibration profile SHALL identify its normative baseline, preserve the core invariants, document every material departure from the global baseline, maintain its required evidence and ratification record, and be versioned so that any historical verdict can be reconstructed. A profile represented as a ratified WEKID sector profile SHALL conform to WEKID-STD-002.

15.4 Claims

An implementation MAY state that it implements or supports WEKID-STD-001 if it meets the applicable technical requirements. Such a statement SHALL NOT be represented as formal WEKID certification unless the implementation has separately completed the WEKID certification or conformance program then in force.

An implementation MAY state that it implements WEKID-STD-001 v1.2.1 including evaluation-independence declaration. An implementation SHALL NOT represent an Independent grade as a WEKID endorsement of its evaluator.

Claim condition	Conforming outcome
Hard gate triggered	AL-0 unless a stricter human reservation governs the decision class.
Soft gate triggered	The capped score resolves under §9.2 and the §10.1 bands; the output SHALL NOT be resolved to AL-0 on the strength of the soft gate alone.

15.5 Transition for Evaluation Independence

For a transition window of two quarters from the publication of Version 1.2, an undeclared Independence Grade SHALL be recorded in the verdict as undeclared rather than defaulting to Shared, and the undeclared-grade condition in Section 9.3 does not apply. After the window, the Shared default in Section 3.7 and the fail-closed condition in Section 9.3 take full effect.

A sector calibration profile ratified before Version 1.2 remains valid. Recalibration is not required solely by this amendment, but the evaluator assurance basis required by Section 14.2 SHALL be added at the profile's next scheduled review under WEKID-STD-002 Provision 10.

© 2025–2026 WEKID LLC. WEKID™ is a trademark of WEKID LLC. | Normative corrective release

Annex A — Baseline Authority Levels (Normative)

This Annex restates the baseline Authority Level definitions. These definitions are normative.

Level	Definition
AL-0 — No Authority	Informational or rejected. No decision, action, or downstream automation may rely on the output.
AL-1 — Constrained	Only whitelisted, reversible, impact-capped actions within pre-approved bounds.
AL-2 — Supervised	AI proposes and stages; an accountable human releases each consequential action before effect.
AL-3 — Augmented	AI executes autonomously within assigned scope, with logging, sampling review, rollback, and human-on-the-loop oversight.
AL-4 — Human Authority	The decision is not delegated. AI is limited to analysis and decision support; an accountable human decides.

Annex B — Baseline Gates and Decision Matrix (Normative)

B.1 Baseline Gates

Gate	Class	Trigger	Effect
Data Integrity Gate	Hard	Data < 2	Cap total score at 50/100; authority resolves under precedence rules.
Fabrication Gate	Hard	Fabricated citation, source, fact, or tool result detected	Cap total score at 40/100; treat as epistemic-integrity violation.
High-Stakes Judgment Gate	Soft	High-stakes content lacks required uncertainty acknowledgment or risk language	Cap total score at 60/100 and require constrained/human review as applicable.
Evaluation Independence Gate	RFC-010 score cap	Data layer graded Shared and no deterministic verification applied to the Data layer	Cap total score at 70/100; Fabrication Gate detection SHALL NOT be relied upon as discharged.

© 2025–2026 WEKID LLC. WEKID™ is a trademark of WEKID LLC. | Normative corrective release

B.2 Baseline Matrix

Apply the fixed precedence and conservative resolution in §9.2. See §10.1 for required actions.

Precedence / condition	Outcome
1. Policy-reserved decision class	AL-4 Human Authority
2. Hard gate or other mandatory rejection condition	AL-0 unless a stricter human reservation governs the decision class
2b. Evaluation Independence Gate (Data graded Shared, no deterministic verification)	Maturity Score capped at 70; resolve remaining conditions against the capped score
3. Overseer-engagement decay under an activated condition	More conservative of AL-2 and one level below the otherwise-resolved level
4. Any required layer < 3, or Wisdom/Experience below profile threshold	AL-2
5. Maturity Score < 50 and no hard gate	AL-0
5. Maturity Score 50–69 or soft gate	AL-1
5. Maturity Score 70–79, no hard or soft gates	AL-3 with enhanced monitoring
5. Maturity Score ≥ 80, no hard or soft gates	AL-3
6. Stakes-floor adjustment	Apply the stakes floor under §7.3

© 2025–2026 WEKID LLC. WEKID™ is a trademark of WEKID LLC. | Normative corrective release

Annex C — Minimum Verdict Contract (Normative)

A machine-readable or human-readable WEKID verdict SHALL expose the following minimum fields or their unambiguous equivalents.

Field	Required content
standard_version	WEKID-STD-001 version used
profile_id / profile_version	Calibration profile and version
evaluation_id	Unique evaluation identifier
evaluated_object	Output, action, or decision-context identifier
timestamp	Evaluation time
evaluator_identity	Evaluator identifier and assurance version
layer_scores	Data, Information, Knowledge, Experience, Wisdom
independence_grades	Per-layer grade — Shared, Attenuated or Independent — and determination basis
maturity_score	Composite score after applicable gating/caps
gates	Triggered gates and their effects
stakes_reservations	Applicable stakes/policy reservations
overseer_engagement	Applicable engagement status/signals when used
authority_level	AL-0, AL-1, AL-2, AL-3, or AL-4
binding_limit	The determinant that produced the assigned Authority Level: maturity_ceiling, stakes_floor, or both where the two coincide.
governing_condition	The §9.2 precedence condition recorded as governing (1, 2, 2b, 3, 4, 5 or 6). REQUIRED where more than one condition was satisfied; RECOMMENDED otherwise.
rationale	Structured or textual explanation of the resolution
enforcement_action	Block, constrain, stage, allow, reserve to human, or equivalent

Where a hard gate or a standing reservation governs, binding_limit takes the value corresponding to the limit that would have bound in the absence of that condition, and governing_condition names the condition that did.

Implementations MAY add fields for evidence provenance, control mappings, remediation, confidence, signatures, or sector-specific obligations.

© 2025–2026 WEKID LLC. WEKID™ is a trademark of WEKID LLC. | Normative corrective release

Annex D — Transition and Provenance (Informative)

WEKID-STD-001 v1.0 is derived from the normative core of the WEKID Master Whitepaper v1.3 (August 2026). The Whitepaper's explanatory history, worked examples, failure cases, solution catalogue, and research discussion are not reproduced here except where necessary to define a requirement.

Ratification record. WEKID-STD-001 v1.0 was reviewed, approved, and released through the WEKID RFC process in September 2026. The approving RFC remains the governance record for establishment of this Standard.

The source Whitepaper integrated two community contributions ratified for Version 1.3: overseer-engagement extensions to Model Two and formalization of the Experience layer. This Standard carries those resulting concepts forward as part of the baseline architecture.

Version 1.2. WEKID-STD-001 v1.2 incorporates WK-CORE-RFC-010, Evaluator Independence and Correlated Epistemic Failure, ratified through the WEKID RFC process in October 2026. The RFC extends to the evaluation path an independence principle the framework already applies to Experience (Section 3.5) and to red-team review (WEKID-STD-002 Phase 8). Companion amendments are carried in WEKID-STD-002 v1.1. RFC-010 remains the governance record for these changes, including its rationale, evidence base, and public-comment disposition.

Version 1.2.1 corrigendum. WEKID-RFC-006 through WEKID-RFC-009 were accepted on October 21, 2026 and omitted from the v1.2 publication. This corrective edition incorporates their operative provisions. Historical verdicts continue to follow the Standard and calibration-profile versions recorded in those verdicts.

The companion research paper Dynamic Epistemic Authority is intentionally not incorporated into this Standard. Its mathematical treatment of maturity trajectory, Trust Bridge admissibility margin, and repeated authority resolution remains non-normative research unless and until adopted through a Standards-Track RFC.

Future normative changes shall proceed through the WEKID-RFC-NNN process and be incorporated into a numbered Standard release.

Source Basis

Judge, James M., and Members of the WEKID Community. WEKID™ Master Whitepaper — A Hierarchical Epistemic Framework and Operational Model for Evaluating and Governing Enterprise Artificial Intelligence Outputs. Version 1.3. WEKID LLC, August 2026.

Judge, James M. WK-CORE-RFC-010: Evaluator Independence and Correlated Epistemic Failure. WEKID LLC, October 2026.

Companion informative research: Judge, James Madigan. Dynamic Epistemic Authority: A Mathematical Model of Maturity, the Trust Bridge, and Decision Authority in the WEKID™ Framework. September 2026.

WEKID-RFC-006 — Conservatism in Precedence Resolution. Council decision October 21, 2026.

WEKID-RFC-007 — Determinate Outcome for Overseer-Engagement Decay. Council decision October 21, 2026.

WEKID-RFC-008 — Gate Class Taxonomy and the High-Stakes Judgment Gate. Council decision October 21, 2026.

WEKID-RFC-009 — Recording the Binding Determinant. Council decision October 21, 2026.

Original v1.2 archive fingerprint (SHA-256)
14a0c9717b3358dd23772d23f5d1656726f3b94a33dbe2fe192af172beb6dd60

© 2025–2026 WEKID LLC. WEKID™ is a trademark of WEKID LLC. | Normative corrective release