AIR-VS 2026.1 — The AIR Vendor Standard

Published by the AIR Standards Council · Version 2026.1 · Effective 2026

The AIR Vendor Standard (AIR-VS) is the public criteria against which AI tools and vendors are assessed for **AIR Certified** status. It is a sibling to the AIR Index (the AI-readiness assessment for organizations adopting AI); both share the AIR brand and standard, but AIR-VS is its own track, with its own surface, scope, and verdict.

Preamble

What this standard attests

AIR-VS attests that an AI tool meets a credible, published trust bar for being placed **near an organization's real work and real data**. It is not a quality ranking, a performance benchmark, or an endorsement of fitness for any specific use. It is an attestation that, on the date of assessment, the vendor's product, contracts, and operating practices clear a defined floor across nine trust dimensions — verified by a human who read the documents and used the product, not inferred from a questionnaire.

Who it's for

Any organization buying or deploying AI to do real work, and any vendor selling an AI tool into any organization, in **any industry**. The dimensions below are deliberately universal: data handling, IP provenance, security, transparency, portability, and oversight are the same concerns whether the buyer is a hospital, a bank, a manufacturer, a public agency, or an independent agency. Independent agencies are where assessment under this standard begins — the first cohort — but nothing in AIR-VS is agency-specific, and the standard is written to be applied identically across verticals.

The core principle

The badge is earned, not bought. The badge is worth exactly what the standard is worth.

A vendor pays a flat, public fee for the **assessment** — for the scrutiny — never for the verdict. The same fee applies whether the vendor passes or fails. There is no recurring listing fee, no license, and no path by which paying more, being larger, or being friendlier changes the outcome. A standard that nobody fails is worthless; AIR-VS is built to fail vendors that should fail, and to publish that it does.

The Dimensions

A vendor is assessed across **nine dimensions**. Each dimension has requirements that contribute to a tier decision **and** an **auto-fail gate** — a non-negotiable floor. A breach of any single gate fails the entire

assessment, regardless of strength elsewhere. Gates are not averaged, traded off, or offset.

Illustrative thresholds below (e.g., specific day-counts, percentages) are flagged as *illustrative* and are fixed precisely in the Council's assessment rubric, which versions alongside this document.

Dimension 1 — Data Handling & Privacy

What it covers. What happens to customer inputs, prompts, documents, and derived data once they enter the product: retention, deletion, sub-processing, workspace/tenant isolation, and — most sharply — whether customer inputs are used to train or improve the vendor's models.

Why a buyer cares. This is the dimension most often misrepresented in marketing and most consequential when wrong. A buyer putting real client data or regulated records into a tool needs to know it will not leak across tenants, will not silently become training data, and can be deleted on request and on exit.

Requirements.

- A written DPA (or equivalent contractual data terms) available to customers, naming retention periods, deletion timelines, and the full sub-processor list.
- Customer inputs are **not** used to train or fine-tune models that serve other customers, unless via explicit, opt-**in**, revocable consent that is off by default. Opt-out-by-default training on inputs does not meet the bar.
- Logical tenant/workspace isolation, verified live in-product (see Process), such that one customer's data and context cannot surface in another's.
- A working deletion path (in-product or contractual SLA) and defined retention defaults.
- Clear handling terms for any third-party model providers the product routes data through (e.g., upstream foundation-model APIs), including whether those providers may retain or train on the data.

> **AUTO-FAIL GATE — Data:** The product trains, fine-tunes, or improves any model on customer inputs **without explicit opt-in consent that is off by default**; OR the vendor cannot produce contractual data terms (DPA or equivalent); OR live verification shows customer data crossing the tenant boundary. Any one is a full assessment failure.

Dimension 2 — IP, Training-Data Provenance & Indemnification

What it covers. The provenance of the data and models the vendor ships, the ownership of outputs, and what the vendor stands behind contractually if its product generates infringing or improperly sourced material.

Why a buyer cares. When a buyer puts AI output into client deliverables, regulated filings, or shipped product, they inherit the vendor's IP exposure. Buyers need to know who owns the output, whether the

underlying models were trained on data the vendor had rights to use, and whether the vendor will stand behind the tool if a claim arises.

Requirements.

- A clear, written statement of **output ownership** assigning rights in generated outputs to the customer (or clearly disclosing any retained vendor rights).
- A disclosed position on **training-data provenance** for any models the vendor itself trains or fine-tunes: licensed, owned, public-domain, or third-party-licensed sources, at a level a buyer can evaluate. "We don't disclose" on owned models does not meet the bar; passing through a third-party foundation model is acceptable if disclosed as such.
- A documented **IP indemnification** position for output-based infringement claims — either an affirmative indemnity, or an explicit, conspicuous disclosure that no indemnity is offered. The bar is honest disclosure, not a mandatory indemnity.
- No contractual terms that claim ownership of customer inputs or customer-derived outputs.

> **AUTO-FAIL GATE — IP:** The vendor's terms claim ownership of customer inputs or customer outputs; OR the vendor materially misrepresents training-data provenance or indemnification (states an indemnity or a provenance claim that the contract or evidence contradicts). Misrepresentation fails; honest "no indemnity offered" does not.

Dimension 3 — Security Posture

What it covers. The vendor's technical and organizational security controls: encryption, access control, vulnerability management, incident response, and independent attestation.

Why a buyer cares. An AI tool with access to real work is an attack surface and a custodian of sensitive data. A buyer needs evidence — not assurances — that the vendor runs a defensible security program and will tell them when something goes wrong.

Requirements.

- Encryption in transit and at rest, documented.
- Role-based access control and a stated internal access policy governing which vendor staff can see customer data and under what conditions.
- A documented vulnerability-management and patching process.
- A written incident-response process including a customer breach-notification commitment with a stated timeline (*illustrative:* notification within 72 hours of confirmed breach).
- Current independent security attestation (e.g., SOC 2 Type II, ISO 27001, or equivalent) for higher-tier certification (see Tiers). For base certification, a documented internal program is required; independent attestation is required for AIR Certified+.

> **AUTO-FAIL GATE — Security:** No encryption of customer data at rest or in transit; OR no incident-response/breach-notification commitment of any kind; OR evidence of an unremediated, known critical vulnerability the vendor declined to address. Any one is a full assessment failure.

Dimension 4 — Transparency, Explainability & AI-Use Disclosure

What it covers. Whether the vendor is honest about what is AI, how it works at a level the buyer can act on, its known limitations, and what the system can and cannot be relied on to do.

Why a buyer cares. Buyers cannot manage risk they cannot see. They need to know which parts of the product are AI-driven, what models or model classes are in use (at a useful level), where the system is known to be unreliable, and what data flows where — so they can scope use, set human checks, and disclose appropriately to their own clients.

Requirements.

- Clear disclosure of which product functions are AI-generated or AI-assisted.
- Disclosure of model provenance at a useful level (own models vs. named third-party providers vs. open-weights), sufficient for a buyer to reason about data flow and dependency.
- A published, honest statement of **known limitations and failure modes** — not marketing, but the real "do not rely on this for X" boundaries.
- Documentation a non-specialist buyer can actually read and act on (not solely an API reference).

> **AUTO-FAIL GATE — Transparency:** The vendor conceals or misrepresents that a core function is AI-driven, OR misrepresents which models/providers process customer data in a way that would mislead a buyer's risk assessment. Active deception about what the system is fails the assessment.

Dimension 5 — Integration & Fit (including Exportability)

What it covers. How the tool connects to a buyer's existing stack, the openness of its interfaces, and — critically — whether a buyer can get their data **out** in a usable form.

Why a buyer cares. A tool that cannot integrate or cannot export traps data and work inside it. Exportability is the practical test of whether a buyer is a customer or a hostage.

Requirements.

- Documented integration surface (APIs, webhooks, or supported connectors) appropriate to the product category, or an honest statement of its absence.
- A working **data-export capability** that returns customer-created data and content in a standard, machine-readable, non-proprietary format (*illustrative:* CSV, JSON, or standard document formats), verified live in-product during assessment.

- Documentation sufficient for a buyer's technical staff to implement integration and export without vendor-gated services.

> **AUTO-FAIL GATE — Export:** The buyer cannot extract their own data and content from the product in any usable, non-proprietary form — i.e., there is no functioning export path. Verified live; an export that exists only in documentation but does not work fails.

Dimension 6 — Portability / No Lock-In

What it covers. What happens to the buyer at the end of the relationship: data return, data destruction, contractual exit terms, and the absence of structural traps that make leaving prohibitive.

Why a buyer cares. Lock-in turns a tool decision into an irreversible dependency. Buyers need confidence that they can leave with their data, that the vendor will destroy what remains, and that exit terms are stated up front — not discovered at renewal.

Requirements.

- Stated **exit terms**: how a customer terminates, what is returned, in what format, and within what window.
- A **data-destruction-on-exit** commitment with a stated timeline.
- No contractual provisions that forfeit customer data or content on termination, and no proprietary-format constructions whose only function is to make exit impractical.
- Reasonable continuity of access to export tooling through a defined wind-down period after termination.

> **AUTO-FAIL GATE — Lock-In:** Contractual terms cause the customer to forfeit or lose access to their own data and content upon termination, OR there is no committed mechanism to return or allow extraction of customer data on exit. Hostage-taking on exit fails.

Dimension 7 — Claims Substantiation

What it covers. Whether the vendor's public and sales claims — about accuracy, security, privacy, performance, certifications, and capabilities — are evidenced and true.

Why a buyer cares. The entire value of a trust badge collapses if vendors can make unsupported claims around it. Buyers rely on a vendor's representations; AIR-VS exists in part to verify that those representations are real.

Requirements.

- Material public claims (accuracy figures, security certifications, privacy guarantees, "we never train on your data," "SOC 2 certified," "bank-grade encryption," etc.) must be substantiated with evidence

on request.
- Certification and compliance claims must correspond to current, verifiable attestations.
- Quantitative performance claims must cite a methodology or evidence base; unfalsifiable superlatives presented as fact are flagged.
- No claim in marketing that the vendor's own contracts or product contradict.

> **AUTO-FAIL GATE — Claims:** A material public claim is demonstrably false or unsubstantiated and the vendor declines to correct it within the cure window — in particular any false claim about data training, security certification, or privacy. A vendor that lies in market and won't fix it fails.

Dimension 8 — Human Oversight & Safety

What it covers. Whether the product is designed so that a human can stay meaningfully in control: oversight points, the ability to review or override AI actions, guardrails on high-impact actions, and safety controls against foreseeable harmful or abusive use.

Why a buyer cares. Buyers are accountable for what the tool does in their name. They need to be able to insert human judgment at the points that matter, override the system, and trust that the vendor has thought about misuse and built brakes — especially as tools take autonomous actions.

Requirements.

- Identified points of **human oversight**: the buyer can review, approve, or override consequential AI outputs or actions where the product takes them.
- For products that take autonomous or high-impact actions, documented **guardrails, confirmation steps, or limits** proportionate to the impact.
- Documented **safety controls** addressing foreseeable misuse and harmful-output categories relevant to the product's domain.
- An accessible mechanism for customers to report safety issues and a stated response commitment.

> **AUTO-FAIL GATE — Oversight:** The product takes consequential, irreversible actions on a buyer's behalf with **no available human-oversight, override, or confirmation mechanism**; OR the vendor has no safety controls of any kind for a product where foreseeable misuse causes material harm. Ungovernable autonomy fails.

Dimension 9 — Reliability & Support

What it covers. Whether the product is operationally dependable and whether the vendor stands behind it: uptime, change management, support responsiveness, and operational maturity.

Why a buyer cares. A tool embedded in real work becomes a dependency. Buyers need to know it will be available, that breaking changes won't arrive unannounced, and that someone will answer when it fails.

Requirements.

- A stated availability/uptime target or an honest statement of service-level posture.
- A change-management practice that gives customers reasonable notice of breaking changes and deprecations.
- A defined support channel with a stated response expectation appropriate to the product tier.
- Evidence of operational maturity (status page, incident history, or equivalent) appropriate to the vendor's stage.

> **AUTO-FAIL GATE — Reliability:** The vendor pushes breaking changes with no notice mechanism **and** offers no support channel of any kind — i.e., a buyer has no way to be warned and no way to get help. Operational abandonment fails.

Scoring & Tiers

AIR-VS uses **gated scoring**. Scoring is sequential, not additive:

1. **Gates first.** A vendor must clear **every** auto-fail gate across all nine dimensions. A single gate breach ends the assessment in a fail, with no offsetting.
2. **Then thresholds.** A vendor that clears all gates is scored against per-dimension thresholds defined in the Council's rubric. The vendor must meet the threshold in **each** dimension — strength in one dimension cannot rescue a deficit in another. Dimensions are floors to clear, not points to accumulate.

Public tiers

- **AIR Certified** — Clears all nine gates and meets every per-dimension threshold, including a documented internal security program.
- **AIR Certified+** — Everything required for AIR Certified, **plus independent third-party attestation** on the security and data dimensions (e.g., current SOC 2 Type II, ISO 27001, or equivalent independent audit). The "+" signals externally verified, not self-attested, controls on the two highest-stakes dimensions.
- **Not Certified** — Did not clear a gate or did not meet a threshold. **This result is private** to the vendor. Individual failures are never published — which is precisely what makes vendors willing to apply.

What is published

The public registry publishes a **tier/band and a date stamp** — *AIR Certified* or *AIR Certified+*, with the standard version and assessment date. **AIR-VS never publishes a raw leaderboard score.** A vendor is in a band or not; there is no ranked number to game.

What is published in aggregate

Each cohort's **aggregate outcome is published**: "Assessed N tools this cohort; M met the bar." Individual failures stay private; the failure rate is public. A badge nobody fails is worthless — the published aggregate is the proof that the bar is real.

The Process

1. **Application + self-attestation.** The vendor applies and completes a structured self-attestation across the nine dimensions. This is the vendor's claim, not the verdict.
2. **Evidence intake.** The vendor submits supporting evidence: DPA/data terms, ToS, security attestations, model and provenance documentation, export/integration docs, and substantiation for material public claims.
3. **Human verification.** A trained assessor verifies, by hand, not by checklist:
4. **Reads the actual DPA and ToS** — line by line — against the self-attestation, looking for contradictions (especially on training, ownership, and exit).
5. **Performs a live, in-product check** of tenant/workspace isolation and the data-export path — confirming export actually works and produces usable, non-proprietary output.
6. **Spot-tests material claims** against evidence — certifications, "never trains on your data," accuracy figures, security representations.
7. **Four-eyes review.** Any **gate decision or boundary case** is reviewed by a second qualified assessor before it stands. No single assessor can issue a fail-on-gate or a borderline pass alone.
8. **Findings + 30-day cure window.** The vendor receives written findings. For curable deficiencies (e.g., an unsubstantiated claim, a missing disclosure, a fixable export gap), the vendor has **30 days** to remediate and resubmit evidence. The cure window does not apply to deliberate misrepresentation.
9. **Decision.** The Council issues the verdict: **AIR Certified**, **AIR Certified+**, or **Not Certified**. The verdict is the Council's, not the firm's.
10. **Public registry entry.** A pass produces a registry entry stamped with the **standard version and assessment date**. The buyer consults the registry independently; certification status lives only there.
11. **12-month re-assessment + material-change clause.** Certification is valid for **12 months**. Expiry triggers a **full re-assessment** — new work, a new fee — not a renewal. Additionally, a **material change** (to data handling, training practices, ownership terms, security posture, or any gated

practice) obligates the vendor to notify the Council and may trigger re-assessment or revocation before the 12 months elapse.

Governance & Integrity

The detail lives in the **AIR Standards Council Charter**; in summary:

- **Independent Council owns the criteria, the verdicts, and revocation.** The AIR Standards Council holds final authority over pass/fail and over revoking a certification. It holds no equity in Nfinit Monkeys and does not depend on certification revenue. **The firm cannot grant a pass or block a revocation.**
- **Fee is labor-based, flat, public, and non-contingent.** The assessment fee is priced to the **labor of assessment**, not the vendor's size or revenue. The **same fee applies whether the vendor passes or fails**, with no refund on a fail. **There is no recurring listing, license, or renewal fee — ever.** Expiry means re-assessment (new work, new fee). The invoice says **"assessment."** You pay for scrutiny, not the verdict.
- **Advisory firewall.** Any AIR Index organizational advisory engagement recommends **no specific vendors** (certified or not) and uses vendor-blind criteria. Certification status lives only on the public registry the buyer consults independently. **No advisory engagement runs concurrently with assessment of the same vendor** (recusal).
- **Aggregate results published.** Each cohort's pass/fail aggregate is public; individual failures are private.
- **The bright line.** The day certification revenue creates pressure to pass a vendor that should fail, the program shuts down. The standard never bends. Neutrality is the franchise; the certification is expendable.
- **Not a warranty.** An AIR-VS assessment is a point-in-time, evidence-based attestation that the vendor cleared the standard on the assessment date. **It is not a warranty, guarantee, or insurance** of the vendor's product, security, or conduct, and it does not transfer the buyer's own due-diligence obligation. The buyer remains responsible for fitness-for-purpose decisions.

Versioning

AIR-VS is versioned as `YEAR.RELEASE` — this is **2026.1**: the first release of the 2026 standard.

- **Release increment** (e.g., 2026.1 → 2026.2): within-year revisions — clarifications, tightened thresholds, new evidence requirements, or rubric refinements that do not redefine the dimensions.
- **Year increment** (e.g., 2026.x → 2027.1): the annual baseline review, in which the Council reviews every dimension and gate against the state of AI tooling, regulation, and observed evasion, and re-baselines the standard.

- **Cadence.** The Council conducts a **scheduled annual review** and may issue **out-of-cycle releases** when a material shift in technology, regulation, or vendor behavior requires it.
- **Stamping.** Every registry entry and every certificate names the **exact standard version** it was assessed against. A certification under 2026.1 is a 2026.1 certification; it does not silently inherit later versions. Re-assessment is performed against the standard version current at the time of re-assessment.
- **Transparency of change.** Each version is published with a changelog describing what changed and why, so buyers and vendors can see how the bar moves over time.

AIR-VS 2026.1 is published and governed by the AIR Standards Council. Certification under this standard is an assessment, not a warranty. Illustrative thresholds noted in-text are fixed in the Council's versioned assessment rubric.