aiBlue Core™

Institutional AI Governance Infrastructure

AIBLUE-INTL-STD-004

AI Governance, Audit and Certification Standard

aiBlue International Standard

Document Code	aiBlue-INTL-STD-004
Classification	Institutional Public Standard
Version	1.0 — 2025
Applicability	International — applicable across jurisdictions and sectors
Intended Audience	Regulators, Auditors, Enterprise Boards, Standards Bodies, Institutional Investors, Certification Bodies, Ecosystem Participants
Relationship	Derived from and consistent with aiBlue-ARCH-003, aiBlue-GOV-001, aiBlue-LIAB-002
Status	Effective upon publication

> *"Governance without auditability is intention without evidence. This Standard translates institutional governance principles into verifiable, auditable, and certifiable operational requirements — making the responsible deployment of artificial intelligence demonstrable across all jurisdictions, not merely declared."*
> *— Foundational Principle, aiBlue-INTL-STD-004*

SECTION 1 | *Version 1.0 — 2025*

Purpose of the Standard

This document establishes the aiBlue AI Governance, Audit and Certification Standard (hereinafter "the Standard" or "aiBlue-INTL-STD-004"). The Standard constitutes a public governance framework applicable to organizations that operate AI-assisted decision environments under the aiBlue Core™ governance infrastructure. It defines the governance baseline, audit protocols, accountability mechanisms, and certification structure against which such organizations may be independently assessed and formally recognized.

The Standard has been developed to serve four complementary institutional functions:

- Governance Baseline: to define the minimum governance requirements against which organizations may be evaluated, certified, and continuously monitored as responsible operators of AI-assisted decision systems;
- Auditability: to establish the audit protocols, traceability requirements, and documentation standards that support independent verification of AI-assisted decision environments by internal auditors, external assurance providers, and regulatory authorities;
- Accountability: to provide a structured, publicly available reference against which regulators, supervisory authorities, institutional investors, and enterprise boards may assess the governance maturity and accountability posture of organizations operating AI systems under the aiBlue governance architecture;
- International Interoperability: to align aiBlue ecosystem governance requirements with the principal international AI governance frameworks, including the NIST AI Risk Management Framework, the OECD AI Principles, the EU AI Act, and the UNESCO Recommendation on the Ethics of AI, such that compliance with this Standard supports — and is recognized as supporting — compliance with those instruments.

> *The aiBlue Core™ governance layer occupies the same institutional role in AI governance that a well-architected cloud governance framework occupies in technology infrastructure management: it does not constitute the AI models or applications themselves, but defines the governance, control, and accountability layer within which those components operate responsibly, traceably, and in alignment with institutional risk standards.*

SECTION 2 | *Version 1.0 — 2025*

Institutional Nature of the Standard

The aiBlue-INTL-STD-004 is an institutional governance instrument. It does not constitute a contractual document, software license, or product specification. It is a public standard that defines the behavioral, procedural, and technical requirements for the responsible governance of AI systems within the aiBlue ecosystem.

2.1 Position within the aiBlue Institutional Architecture

This Standard occupies a specific tier within the aiBlue institutional documentation hierarchy. It formalizes and operationalizes the governance content present across the ecosystem's foundational instruments:

Instrument	Nature	Relationship to This Standard
aiBlue-ARCH-003	Architecture Declaration	Defines the three-layer technical and institutional structure within which this Standard's requirements operate.
aiBlue-GOV-001	AI Governance Charter	Establishes governing principles and internal governance architecture. This Standard operationalizes those principles into auditable requirements.
aiBlue-LIAB-002	Responsibility Framework	Delineates actor responsibility allocation. This Standard formalizes those allocations into audit-ready accountability structures.
aiBlue-INTL-STD-004	Governance Standard (this document)	Public certification and audit framework derived from and consistent with all prior instruments.

2.2 Adoption and Binding Effect

Compliance with this Standard is:

- Mandatory for all organizations seeking aiBlue Governance Certification under the Programme defined in Section 9;
- Incorporated by reference into platform participation agreements between aiBlue and its ecosystem participants, to the extent specified in such agreements;
- A publicly available reference document accessible to regulators, supervisory authorities, certification bodies, auditors, and institutional counterparties worldwide.

SECTION 3 | *Version 1.0 — 2025*

Scope of Application

This Standard applies to all participants in the aiBlue Core™ governance ecosystem and to all AI-assisted decision environments orchestrated through the aiBlue governance infrastructure, without restriction to any particular jurisdiction, industry sector, or organizational type.

3.1 Organizational Scope

- aiBlue (as proprietor and operator of the governance infrastructure layer);
- Authorized Deployment Partners (ecosystem operators who deploy, implement, and support the Platform for client organizations);
- Client Organizations (enterprises, government bodies, regulated institutions, public authorities, academic institutions, and any other entity that operates AI systems within the aiBlue governance infrastructure).

3.2 Functional Scope

The Standard governs:

- All AI-assisted decision environments in which the aiBlue Core™ governance layer is deployed;
- All AI models — regardless of origin, architecture, or access modality — that are orchestrated, governed, or monitored through the aiBlue governance layer;
- All human actors participating in AI-assisted decision processes under the aiBlue governance architecture, including designated supervisors, AI governance officers, compliance personnel, and organizational decision-makers.

3.3 Exclusions

The Standard does not govern:

- The intrinsic architecture, training methodology, or model behavior of third-party AI systems, which remain the responsibility of their respective developers under applicable law and their own terms of service;
- Business decisions by client organizations that are independent of any AI-assisted process;
- Internal processes of Deployment Partners that do not involve the Platform or client governance operations.

3.4 Jurisdictional Applicability

This Standard is designed to be jurisdiction-neutral in its governance principles, while accommodating sector-specific and jurisdiction-specific compliance requirements through documented overlay procedures. Organizations operating in regulated sectors are expected to identify and apply the specific regulatory obligations of their jurisdiction as supplementary requirements atop the governance baseline established herein. The table below illustrates representative regulatory alignment:

Jurisdiction / Framework	Relevant Instruments	Relationship to This Standard
European Union	EU AI Act; GDPR; NIS2 Directive	High-risk system controls (Sections 6, 7); transparency (Section 5); conformity assessment alignment (Section 9)
United States	NIST AI RMF; Executive Order on AI; CCPA	GOVERN / MAP / MEASURE / MANAGE function mapping; risk classification alignment

United Kingdom	UK AI Safety Framework; UK GDPR; ICO Guidance	Accountability principles; transparency and explainability requirements
Brazil	LGPD; PL 2338/2023; BCB Resolution 85/2021	Right of review provisions; impact assessment requirements; sector-specific controls
International	OECD AI Principles; UNESCO AI Ethics; ISO/IEC standards	Foundational principles alignment (Section 13)

SECTION 4 | *Version 1.0 — 2025*

Governance Architecture Principles

The governance architecture of the aiBlue ecosystem is founded on eight inviolable principles. These principles govern the design, configuration, and operation of all AI-assisted decision environments within the ecosystem and are non-negotiable in their application. They prevail over operational convenience, commercial considerations, and efficiency arguments.

I. Human Centrality

AI systems exist to augment human capabilities, not to displace human responsibility or judgment. Every decision of significant consequence for individuals, communities, or institutions requires structured human supervision, validation, and accountability. No governance configuration under this Standard may be designed to systematically eliminate human oversight from consequential decision processes.

II. Transparency and Explainability

Organizations have a duty to understand, at a level commensurate with the risk of the system, how their AI systems arrive at particular outputs or recommendations. AI-assisted decision processes must be documented in a manner that permits explanation, in accessible language, of the principal factors that influenced each decision — particularly where such decisions affect the rights or interests of individuals.

III. Non-Discrimination and Algorithmic Equity

Organizations must actively prevent, identify, and remediate algorithmic biases that produce discriminatory outcomes on the basis of characteristics protected under applicable law and international human rights standards.

IV. Proportionality and Risk Minimisation

Governance controls must be proportionate to the level of risk associated with each AI use case. Higher-impact, higher-risk applications require more robust controls, more intensive supervision, and more detailed documentation. The risk classification framework defined in Section 6 is the reference instrument for this proportionality determination.

V. Accountability and Answerability

For every AI system in production, a responsible party — an identified person or organizational function — must be designated and capable of answering for the system's decisions, limitations, and impacts. Absence of an identified responsible party precludes approval for production use.

VI. Privacy and Data Protection by Design

Personal data protection is embedded from the design stage of any AI use case. The least privacy-invasive technical option is adopted at every design decision point, absent documented and proportionate justification.

VII. Security and Resilience

AI systems are treated as critical information assets. Cybersecurity and operational resilience measures proportionate to their risk profile are mandatory, including access control, continuous monitoring, vulnerability management, and business continuity planning.

VIII. Continuous Improvement and Regulatory Adaptability

AI governance is not static. Organizations must review their processes, controls, and policies at least annually in light of technological, regulatory, and ethical developments, and must adapt proactively — without awaiting express regulatory compulsion.

4.1 Three-Layer Governance Architecture

The aiBlue governance architecture is structured in three interdependent functional layers. All certification and audit activities under this Standard are conducted against this architectural model:

Layer	Components	Accountable Party
Strategic	AI governance policy; board-level AI risk oversight; high-risk use case approval; enterprise AI strategy	Board of Directors / C-Suite / AI Governance Committee
Tactical	Data Protection Impact Assessments; risk evaluations; compliance programme; internal audit; algorithmic equity monitoring	DPO / CRO / CISO / Compliance function
Operational	Technical controls; continuous monitoring; incident management; audit trails; aiBlue Core™ infrastructure	Technology / Engineering / AI Governance Officer

SECTION 5 | *Version 1.0 — 2025*

Operational Transparency Requirements

Operational transparency is a prerequisite for governance legitimacy. This Section defines the minimum transparency requirements applicable to each AI system in production within the aiBlue governance architecture.

5.1 System Transparency Register

For each AI system in production that affects decisions concerning natural persons, the responsible organization must maintain a System Transparency Register containing, at minimum:

- System identification: unique identifier, version number, deployment date, and designated responsible parties (technical owner and business owner);
- Purpose statement: documented description of authorized use cases and the populations or processes to which the system is applied;
- AI model identification: identification of the model(s) in use, their provenance (third-party provider or internal development), and the deployed version;
- Input variables: description of the categories of data inputs the system processes in generating its outputs or recommendations;
- Performance indicators: the metrics by which system performance is monitored, including accuracy benchmarks, equity indicators, and degradation thresholds;
- Known limitations: documented description of populations, contexts, or conditions in which the system exhibits degraded performance or elevated error rates;
- Oversight level: the designated supervision level (Level A, B, or C as defined in Section 6) and the governance rationale for that classification;
- Explainability mechanisms: description of the technical mechanisms available to produce intelligible explanations of individual outputs or recommendations;
- Applicable regulatory instruments: identification of the regulatory and legal framework governing this system's deployment context.

5.2 Public Transparency Disclosure

Organizations operating at Certification Level 2 or above must publish, on a publicly accessible channel, a summary disclosure for each high-impact AI system covering: its general purpose; the categories of decisions it influences; the level of human supervision applied; any right of individual review or contestation available under applicable law; and a contact point for inquiries and complaints.

5.3 Internal Governance Reporting

Governance reporting to internal oversight bodies must include quarterly reporting to the AI Governance Committee covering system performance indicators, anomaly alerts, equity metrics, incident logs, and regulatory compliance status; and an annual board-level report covering consolidated AI risk exposure, certification status, regulatory developments, and an evaluation of the organization's overall AI governance maturity.

SECTION 6 | *Version 1.0 — 2025*

Human Oversight Protocols

Human oversight is the cornerstone of responsible AI governance under this Standard. Every AI system deployed within the aiBlue governance architecture must be assigned a supervision level before approval for production operation. The supervision framework is risk-based: governance controls scale proportionally to the potential impact of each AI use case.

6.1 Mandatory Supervision Level Classification

Level	Definition	Minimum Requirements
Level A — Monitored Autonomous	Decisions of minimal impact, reversible in character, not directly affecting natural persons as their subject, with pre-defined criteria and automated monitoring. Examples: internal document classification, operational queue management, content filtering.	Automated monitoring with anomaly alerts; complete audit log; simplified impact assessment; semi-annual review; AI Governance Officer validation
Level B — Supervised Assisted	AI produces recommendations; a qualified human professional reviews and decides. The AI system cannot execute consequential actions without documented human authorisation. Examples: credit risk scoring, candidate screening, insurance claims assessment.	Full impact assessment; AI Governance Officer approval; human review documented in audit trail; monthly compliance report; right of individual review where required by applicable law; annual governance committee review
Level C — Human-Authorised with AI Support	AI provides analytical inputs only. The decision, documentation, and accountability are entirely human. The AI system may not generate binding recommendations. Examples: benefit determinations, diagnostic conclusions, disciplinary decisions, decisions materially affecting fundamental rights.	Full impact assessment; governance committee approval; documented right of review guaranteed to affected individuals; annual external audit; prohibition on binding AI outputs; complete human accountability trail

6.2 Prohibited Use Categories

The following categories of AI deployment are prohibited within the aiBlue governance architecture, consistent with applicable international regulatory restrictions and the prohibition categories defined in the aiBlue Responsibility Framework (aiBlue-LIAB-002):

- General social scoring of natural persons based on social behaviour or personal characteristics, where such scoring causes unjustifiable harm or discrimination;
- Subliminal or manipulative techniques that exploit psychological vulnerabilities to distort behaviour or decision-making;
- Indiscriminate biometric mass surveillance in publicly accessible spaces, except where expressly authorised by applicable law for specific, proportionate purposes;
- AI-based profiling that systematically produces outcomes incompatible with the non-discrimination principle established in Principle III of Section 4.

6.3 Governance Structure Requirements

Organizations must establish and maintain: an AI Governance Committee with qualified cross-functional membership meeting at least quarterly; a designated AI Governance Officer with authority over use case

approval, monitoring oversight, and incident escalation; and documented procedures for individual review and contestation rights, including response timelines, responsible persons, and complete audit logging of review requests and outcomes.

SECTION 7 | *Version 1.0 — 2025*

Auditability and Decision Traceability

Auditability is the technical and procedural capacity to reconstruct, verify, and account for AI-assisted decisions after the fact. It is the primary mechanism through which governance principles are made demonstrable rather than declaratory. The aiBlue Core™ infrastructure provides the technical foundation for auditability; this Section defines the governance requirements that must be met by organizations operating within that infrastructure.

7.1 Mandatory Audit Trail Components

For each AI-assisted decision of consequential impact, the audit trail must capture and preserve:

- Decision event timestamp: precise date, time, and system context of the AI operation;
- System and model identification: unique identifier of the AI system and the specific model version in use at the time of the decision;
- Input data record: the categories and, where legally permissible, the specific data inputs provided to the AI system for the particular decision;
- AI system output: the recommendation, classification, score, prediction, or other output generated by the system;
- Human action record: the identity and role of the human actor who reviewed the AI output (at Level B and Level C), the decision taken, and documented rationale for any departure from the AI recommendation;
- Communication record: timestamp and channel through which the final decision was communicated to any affected party;
- Review and contestation record: log of any requests for review or contestation received in connection with the decision, and the organisational responses provided.

7.2 Immutability and Integrity Standards

Audit trail records must be stored in an immutable format with cryptographic integrity controls that make post-generation modification detectable; preserved for a minimum of five years, or such longer period as required by applicable law; and accessible for structured export in formats compatible with governance, risk, and compliance tooling used by the organisation, its auditors, and applicable regulatory authorities.

7.3 Explainability Documentation

For any AI-assisted decision subject to an individual's right of review or explanation under applicable law, the organisation must be capable of producing, upon request, an explanation that: identifies the principal factors influencing the AI system's output for the specific case; is expressed in language accessible to the affected individual or their representative; and provides sufficient information to enable the individual to understand the basis of the decision and exercise any applicable legal rights — without requiring disclosure of proprietary model architecture or trade secrets.

7.4 Retrospective Reconstruction Capability

Organizations must maintain the technical and procedural capability to retrospectively reconstruct, for any past consequential AI-assisted decision, the state of the AI system at the time of the decision (model version, configuration, operating parameters), the inputs processed, the output generated, and the human actions taken in response. This capability is the evidentiary foundation for regulatory investigations, litigation support, internal incident analysis, and independent audit.

SECTION 8 | *Version 1.0 — 2025*

Responsibility Allocation Across Actors

The aiBlue ecosystem involves four categories of actors with distinct and non-interchangeable roles. Clear delineation of responsibility across these actors is a prerequisite for effective governance, proportionate liability management, and credible regulatory accountability. This Section establishes the authoritative responsibility allocation framework applicable throughout the ecosystem.

8.1 Actor Architecture

Actor	Governance Role and Primary Accountability
aiBlue	Owner and operator of the governance infrastructure layer. Responsible for the availability, security, integrity, and correct operation of the aiBlue Core™ infrastructure, audit trail generation, and the operation of this Standard and the Certification Programme.
Third-Party AI Model Providers	Developers, trainers, and maintainers of AI models orchestrated through the Platform. Responsible for the intrinsic characteristics of their models, including training data, model behaviour, and output quality, under their own agreements with client organisations.
Client Organisation	Deployer and operator of AI use cases within the governance infrastructure. Responsible for governance configuration, human oversight implementation, regulatory compliance in its operating jurisdiction, and ultimate accountability for consequential decisions made using AI-assisted systems.
Human Decision-Makers	Qualified professionals designated to exercise oversight, review AI recommendations, and take accountable decisions within the governance architecture. Responsible individually, and as agents of the client organisation, for the quality of human oversight exercised.

8.2 aiBlue Governance Infrastructure Responsibilities

aiBlue is responsible for: the availability, security, and integrity of the aiBlue Core™ infrastructure in accordance with contractually agreed service parameters; the accuracy and immutability of Platform-generated audit trail records; the correct operation of governance control mechanisms as configured by client organisations; the security of data processed through the Platform in accordance with applicable law; the maintenance and evolution of this Standard and associated governance documentation; and the operation of the Certification Programme defined in Section 9.

> *The responsibility of aiBlue is analogous to that of a provider of high-assurance enterprise infrastructure: aiBlue is accountable for the integrity, availability, and governance capability of the system — not for the content of information processed therein, nor for the business decisions made on the basis of that information.*

8.3 Client Organisation Responsibilities

Client organisations are responsible for: determining, approving, and supervising all AI use cases deployed through the Platform; defining governance policies configured within the Platform, including supervision levels, access controls, and audit retention parameters; ensuring that authorised users are trained and aware of their oversight obligations; continuous monitoring of AI system performance in production; compliance with all regulatory obligations applicable to the organisation's sector and jurisdiction; and all damages caused to third parties by AI-assisted decisions taken within its processes, irrespective of which AI model generated the underlying output.

8.4 Incident Attribution Sequence

In the event of an incident involving an AI-assisted decision, responsibility attribution follows this analytical sequence:

- Infrastructure investigation: whether the incident originated in a failure of the aiBlue Core™ infrastructure (attributable to aiBlue);
- AI model investigation: whether the incident originated in a failure or inappropriate behaviour of a third-party AI model (attributable to the model provider, with potential client co-liability);
- Governance investigation: whether the client organisation had configured adequate controls and designated sufficient human supervision (attributable to the client for governance failure);
- Human decision investigation: whether the designated human actor exercised the required supervision with the diligence expected of a competent professional in their field (attributable to the individual and to the client as employer).

Complete and intact audit trails generated by the aiBlue Core™ infrastructure constitute the primary evidentiary instrument for conducting this attribution analysis.

SECTION 9 | *Version 1.0 — 2025*

Certification Framework

The aiBlue Governance Certification Programme provides a structured, tiered mechanism for organisations to formally demonstrate their compliance with this Standard. Certification constitutes recognition of governance maturity and serves as evidence of responsible AI operation for the purposes of regulatory interactions, board-level reporting, investor due diligence, institutional procurement, and international counterparty assessment.

9.1 Certification Tiers

Level	Designation	Minimum Requirements
Level 1	aiBlue Governance Foundation	aiBlue Core™ deployed; System Transparency Register operational; supervision level classification complete; audit trail generation active; AI Governance Committee established; internal self-assessment completed
Level 2	aiBlue Governance Standard	All Level 1 requirements; public transparency disclosures published; data protection impact assessments completed for applicable systems; annual internal audit conducted; explainability mechanisms tested and documented; incident response plan tested
Level 3	aiBlue Governance Excellence	All Level 2 requirements; annual independent external audit conducted; algorithmic equity assessments documented and remediated; demonstrated alignment with a recognised international framework (NIST AI RMF or equivalent); advanced board-level AI risk reporting; participation in the aiBlue Governance Community of Practice

9.2 Certification Process

9.2.1 Self-Assessment

Organisations initiate certification by completing the aiBlue Governance Self-Assessment Questionnaire, which maps current governance practices against the requirements of the target certification level and identifies gaps to be addressed prior to formal assessment.

9.2.2 Documentation Review

aiBlue or an Accredited Certification Partner conducts a structured review of governance documentation, including the System Transparency Register, impact assessment records, audit trail samples, incident logs, and AI Governance Committee records.

9.2.3 Technical Control Verification

For Level 2 and Level 3 certifications, a technical review of the aiBlue Core™ deployment configuration verifies that governance controls are correctly implemented and that audit trail generation meets the requirements of Section 7.

9.2.4 Certification Issuance

Upon satisfactory completion of the assessment process, aiBlue issues a certification statement specifying the level attained, the scope, the date of issuance, and the validity period of twelve months. Certifications are listed in the public Certification Register maintained by aiBlue.

9.2.5 Accredited Certification Partners

Authorised Deployment Partners with demonstrated governance expertise may apply for accreditation to conduct Level 1 and Level 2 assessments. Level 3 certifications require direct aiBlue involvement or engagement of an externally accredited audit firm. Accreditation criteria are specified in the aiBlue Partner Accreditation Standard.

SECTION 10 | *Version 1.0 — 2025*

Audit and Compliance Procedures

10.1 Continuous Compliance Monitoring

The aiBlue Core™ infrastructure provides automated continuous compliance monitoring, including: real-time performance monitoring with automated alerts for degradation and anomalies; automated audit trail generation for all consequential AI-assisted interactions; data drift and concept drift detection with configurable alert thresholds; algorithmic equity monitoring with segmented metrics by relevant population groups; and access control and user action logging across all interactions with the governance infrastructure.

10.2 Internal Audit Programme

Organisations at Certification Level 2 and above must maintain an annual internal audit programme covering: completeness and accuracy of the System Transparency Register; continued appropriateness of supervision level classifications; operational testing of explainability mechanisms; audit trail record integrity verification; adequacy of human oversight practices; incident response capability and management; and reporting of findings with remediation recommendations to the AI Governance Committee.

10.3 External Audit

Organisations at Certification Level 3 must submit annually to an independent external audit covering all elements of the internal audit programme, plus: independent verification of public transparency disclosures; full governance maturity assessment against this Standard; benchmarking against applicable international frameworks; and a formal audit opinion suitable for inclusion in regulatory submissions, board reporting, and investor disclosures.

10.4 Non-Compliance and Remediation

Material non-compliance identified through audit activities triggers a formal remediation process: written notification to the responsible party with a description of the deficiency; root cause analysis within thirty days; a documented remediation plan with actions, responsible persons, and completion timelines submitted to the AI Governance Committee; verification through a targeted audit review prior to closure; and, where remediation is not achieved within the agreed timeline, potential suspension or downgrade of certification status.

SECTION 11 | *Version 1.0 — 2025*

Evolution and Versioning of the Standard

AI governance is a dynamic discipline. The regulatory landscape, the capabilities of AI systems, and the practices of responsible governance evolve continuously. This Standard is maintained as a living document subject to a structured versioning process designed to ensure that governance requirements remain current, proportionate, and aligned with best international practice.

11.1 Versioning Policy

- Major versions (e.g., v1.0 to v2.0): substantive changes to governance requirements, certification structure, or foundational principles. Requires a minimum of ninety days' prior notice and a structured consultation process with ecosystem participants.
- Minor versions (e.g., v1.0 to v1.1): clarifications, supplementary guidance, or non-substantive additions. Requires thirty days' prior notice.
- Corrigenda: technical corrections of errors or inconsistencies, published immediately with annotation.

11.2 Triggers for Review

Formal review of this Standard is triggered by: annual scheduled review; enactment or substantive amendment of AI-specific legislation in a major jurisdiction; emergence of AI capabilities or deployment patterns creating material governance gaps; critical incidents revealing systematic deficiencies; or formal guidance from regulatory or international bodies affecting the interpretation of current requirements.

11.3 Consultation and Transition

Major version changes are preceded by structured consultation with Deployment Partners, representative client organisations, and relevant regulatory bodies. Upon publication of any new version, organisations with existing certifications receive a grace period equal to the applicable notice period to achieve compliance. All certifications in the Certification Register note the version of the Standard under which they were issued.

SECTION 12 | *Version 1.0 — 2025*

Relationship with Existing Governance Instruments

This Standard is derived from and consistent with the existing aiBlue institutional documentation suite. It does not supersede or amend any prior instrument; it formalizes the governance content present across those instruments into a public, auditable, and certifiable framework. The following matrix documents the relationship between this Standard and each foundational instrument.

Source Instrument	Content Formalised in This Standard	Sections
aiBlue-ARCH-003 Architecture Declaration	Three-layer architecture; institutional positioning of aiBlue Core™ as governance and control layer distinct from AI model providers; regulatory compliance basis	2, 3, 4, 8
aiBlue-GOV-001 AI Governance Charter	Eight foundational governance principles; three-layer governance architecture; governance committee structure; supervision level classification; AI model lifecycle governance; transparency and explainability obligations; algorithmic risk monitoring; incident classification	4, 5, 6, 7, 10
aiBlue-LIAB-002 Responsibility Framework	Actor architecture and responsibility delineation; risk classification framework; decision documentation requirements; incident attribution sequence; regulatory alignment	3, 8, 10
Platform Participation Agreements	Ecosystem structure; intellectual property boundaries; AI model separation provisions; compliance verification rights; data treatment obligations; governance administration	2, 3, 9

12.1 Interpretation Principles

Where any ambiguity or apparent inconsistency exists between this Standard and any prior aiBlue institutional instrument, the following interpretation principles apply: this Standard does not modify or override any contractual obligation in platform participation agreements; governance principles in aiBlue-GOV-001 are the primary reference for interpretation; responsibility allocations in aiBlue-LIAB-002 govern accountability attribution; and architectural constraints in aiBlue-ARCH-003 define the technical implementation boundaries within which controls must operate.

SECTION 13 | *Version 1.0 — 2025*

Alignment with International Governance Frameworks

This Standard has been designed to be structurally compatible with the principal international AI governance frameworks currently in force or in advanced development. Organisations certified under this Standard thereby obtain a structured basis for demonstrating alignment with those frameworks. The following matrix documents the principal correspondences.

International Framework	Principal Correspondences in This Standard
NIST AI Risk Management Framework (AI RMF 1.0, 2023)	GOVERN function: Section 4 (principles), Section 6 (oversight protocols), Section 2 (governance structure). MAP function: Section 3 (scope), Section 6.1 (risk classification). MEASURE function: Section 7 (auditability), Section 5 (transparency). MANAGE function: Section 10 (audit and compliance), Section 9 (certification).
EU AI Act (Regulation 2024/1689)	Prohibited practices: Section 6.2. High-risk system requirements: Sections 5, 6, 7. Transparency obligations: Section 5. Human oversight requirements: Section 6. Technical documentation: Section 7. Conformity assessment: Section 9. Post-market monitoring: Section 10.1.
OECD AI Principles (2019, updated 2024)	Inclusive growth, sustainable development, and well-being: Principle IV. Human-centred values and fairness: Principles I, III. Transparency and explainability: Principle II. Robustness, security, and safety: Principle VII. Accountability: Principle V.
UNESCO Recommendation on the Ethics of AI (2021)	Human rights and human dignity: Principles I, III. Transparency and explainability: Principle II. Responsibility and accountability: Principle V. Privacy and data protection: Principle VI. Sustainability: Principle VIII.
ISO/IEC Standards (AI Management System; AI Risk Management)	Management system approach: Section 4.1. Risk classification: Section 6.1. Audit and review: Section 10. Continual improvement: Section 11. Documentation requirements: Sections 5, 7.
Major Data Protection Frameworks (GDPR; LGPD; CCPA; PIPL and equivalents)	Right of review and explanation for automated decisions: Sections 5.2, 6.3, 7.3. Impact assessment requirements: Sections 6.1, 9.1. Data minimisation and purpose limitation: Principle VI. Incident notification: Section 10.4. Controller accountability: Section 8.3.

Normative References

[01] NIST AI Risk Management Framework (AI RMF 1.0), National Institute of Standards and Technology, 2023
[02] OECD Recommendation on Artificial Intelligence, OECD/LEGAL/0449, 2019 (updated 2024)
[03] UNESCO Recommendation on the Ethics of Artificial Intelligence, 2021
[04] Regulation (EU) 2024/1689 — EU Artificial Intelligence Act

[05] Regulation (EU) 2016/679 — General Data Protection Regulation (GDPR)
[06] ISO/IEC 42001:2023 — Artificial Intelligence Management System
[07] ISO/IEC 23894:2023 — Guidance on Risk Management for AI
[08] ISO/IEC 38500:2024 — Governance of IT for the Organisation
[09] IEEE 7000-2021 — Model Process for Addressing Ethical Concerns in System Design
[10] aiBlue-ARCH-003 — Institutional Architecture Declaration
[11] aiBlue-GOV-001 — AI Governance Charter
[12] aiBlue-LIAB-002 — AI Responsibility Framework

Issued by:
aiBlue Technologies
Institutional Governance Division

Document Status:
Effective upon publication
Version 1.0 — 2025 | International Standard