AI Insurance Readiness Score
Open Standard Specification

AIRS v1.0

Document Version 1.0 | April 2026

Standard Version: 1.0 | Initial Publication

Published by

AI Security Intelligence LLC

Persistent URL
https://www.aisecurityintelligence.com/airs-v1-specification

ABSTRACT

The AI Insurance Readiness Score (AIRS) is an open standard specification for quantifying artificial intelligence security risk in entities subject to insurance underwriting. AIRS defines five weighted scoring domains encompassing twenty-five discrete risk factors, a composite scoring methodology producing scores on a 0–100 scale, and four rating tiers that map directly to underwriting decisions. This specification provides the complete methodology, scoring rubrics, rating tier definitions, assessment process requirements, and conformance criteria necessary for institutional adoption by carriers, reinsurers, regulators, and enterprise risk teams.

OPEN STANDARD

Publication Details

Standard Identifier: AIRS v1.0

Document Version: 1.0

Publication Date: April 2026

Standard Class: Open Standard Specification

Publisher: AI Security Intelligence LLC

Jurisdiction: United States of America

How to Cite This Standard

> AI Security Intelligence LLC. (2026). AI Insurance Readiness Score Open Standard Specification (AIRS v1.0, Document Version 1.0). AI Security Intelligence LLC. https://www.aisecurityintelligence.com/airs-v1-specification

Short in-text reference: (AIRS v1.0) or AIRS v1.0

Copyright and Permitted Use

Attribution Requirements

Any organization publishing AIRS scores, referencing AIRS tiers in policy language, or citing AIRS methodology in regulatory filings shall include the following attribution: "Scored using the AI Insurance Readiness Score (AIRS) v1.0 methodology, published by AI Security Intelligence LLC." The standard version shall be specified in all published scores.

Contact

For comments, errata reports, or implementation inquiries:
standards@aisecurityintelligence.com
https://www.aisecurityintelligence.com/contact

Version History

Date	Document Version	Standard Version	Change Type	Description
April 2026	1.0	1.0	Initial Publication	First release of AIRS v1.0 Open Standard Specification

Version numbering convention: AIRS employs a dual version system following CVSS v4.0 practice. The standard version (e.g., v1.0) changes only with substantive methodology revisions that alter scoring outcomes. The document version (e.g., 1.0, 1.1) tracks editorial corrections and clarifications to the specification text that do not change the scoring methodology. This separation ensures that minor textual corrections do not require re-versioning of the standard itself.

Abstract

The AI Insurance Readiness Score (AIRS) is an open standard specification that establishes a comprehensive, quantitative methodology for assessing artificial intelligence security risk in entities subject to insurance underwriting. AIRS addresses a critical gap in the global insurance industry: the absence of a standardized, transparent framework for evaluating AI-specific risk exposures that carriers, reinsurers, and regulators can uniformly adopt.

The standard defines five weighted scoring domains—Model Integrity and Validation, Output Liability and Content Risk, AI Supply Chain Security, Regulatory and Compliance Alignment, and Systemic Resilience and Continuity—encompassing twenty-five discrete risk factors. Each factor is assessed on a five-level maturity scale with rubric-level specificity. Domain scores are combined through a weighted composite formula producing a 0–100 AIRS score, which maps to four rating tiers with direct underwriting implications.

This specification provides the complete methodology necessary for institutional adoption: domain definitions, factor descriptions, scoring rubrics, the composite calculation formula with mathematical notation, rating tier definitions with threshold ranges, assessment process requirements, conformance criteria, and cross-framework alignment tables mapping AIRS to NIST AI RMF 1.0, ISO/IEC 42001:2023, and the EU Artificial Intelligence Act.

Keywords

AI risk, AI risk scoring, AI security, artificial intelligence, composite score, cyber insurance, cybersecurity rating, enterprise risk management, insurance underwriting, maturity model, open standard, quantitative risk assessment, rating methodology, regulatory compliance, reinsurance, risk scoring, scoring rubric, supply chain security, underwriting framework

Foreword

The global insurance industry underwrites risk across virtually every domain of human activity. For centuries, the industry has developed and refined methodologies for quantifying exposures that were once considered unmeasurable—from maritime commerce to cyber liability. Each new category of risk has demanded new frameworks for measurement, new vocabularies for communication, and new standards for institutional consensus.

Artificial intelligence represents the next such category. AI systems are now embedded in critical infrastructure, healthcare delivery, financial services, and national security. The risk exposures they create—from model failure and adversarial attack to supply chain concentration and regulatory non-compliance—are novel, interconnected, and rapidly evolving. Yet the insurance industry currently lacks a standardized methodology for evaluating these exposures in the entities it underwrites.

Existing cybersecurity rating platforms evaluate traditional IT hygiene—patching cadence, email configuration, network exposure—through proprietary, opaque methodologies. While these tools provide useful signals, they were not designed to assess AI-native risks: training data provenance, adversarial robustness, output liability, model supply chain dependencies, or systemic concentration. The gap between what these tools measure and what underwriters need to evaluate grows wider with each passing quarter.

AIRS was developed to close this gap—not as a proprietary product, but as an open standard that the entire industry can adopt, reference, and build upon. The methodology published in this specification defines what to measure and how to score it. It is designed to be cited in policy language, referenced in regulatory filings, and incorporated into reinsurance treaty terms. A standard is not a tool. A tool solves a problem. A standard defines how an industry agrees to measure.

This specification was developed by AI Security Intelligence LLC through analysis of over 500 AI security incidents, regulatory framework review across 40+ jurisdictions, and cross-reference with established standards including NIST AI RMF 1.0, ISO/IEC 42001:2023, the EU AI Act, CVSS v4.0, and FAIR v3.0. The framework is informed by actuarial principles and calibrated to the decision architecture of institutional underwriting.

Table of Contents

Section 1: Introduction

1.1 Background and Context

The integration of artificial intelligence into enterprise operations has created a new category of risk exposure that existing insurance frameworks were not designed to address. Traditional cybersecurity rating methodologies evaluate perimeter-level hygiene—patching frequency, email authentication, network configuration—through outside-in scanning. These signals remain relevant but are insufficient for quantifying AI-specific risks including model integrity failures, adversarial vulnerabilities, output liability, supply chain dependency concentration, and regulatory non-compliance across evolving jurisdictional requirements.

The global cyber insurance market, projected to exceed $20 billion in gross written premium by 2027, increasingly confronts AI-related claims for which no standardized underwriting methodology exists. Carriers currently rely on self-reported questionnaires, point-in-time penetration tests, or proprietary vendor scores that lack transparency, reproducibility, and AI-native risk coverage. This absence creates pricing uncertainty, adverse selection, and regulatory exposure for carriers, while leaving enterprises without a clear roadmap for demonstrating AI security maturity to their insurers.

AIRS was developed to address this structural gap by establishing an open, transparent, and reproducible methodology for AI security risk scoring. Unlike proprietary rating products, AIRS publishes the complete methodology—domains, factors, weights, scoring rubrics, and tier definitions—as a freely referenceable standard that carriers, reinsurers, regulators, and enterprises can adopt independently of any single vendor.

1.2 Purpose and Objectives

This specification has four primary objectives:

- Define the scoring methodology. Establish the complete AIRS framework including five scoring domains, twenty-five risk factors, a five-level maturity rubric, weighted composite scoring, and four rating tiers.
- Enable institutional adoption. Provide sufficient methodological detail for carriers, reinsurers, and regulators to implement AIRS-conformant assessments within existing underwriting workflows.
- Establish conformance criteria. Define what constitutes conformant use of the AIRS standard, including mandatory requirements, permitted adaptations, and prohibited modifications.
- Facilitate cross-framework alignment. Map AIRS to established regulatory frameworks including NIST AI RMF 1.0, ISO/IEC 42001:2023, and the EU AI Act, enabling organizations to leverage existing compliance investments.

1.3 Scope and Applicability

AIRS is designed to assess AI security risk in entities that insurers underwrite—that is, organizations deploying, operating, or integrating AI systems as part of their business operations. The standard evaluates outward-facing AI risk posture, not internal insurer AI governance.

> Scope distinction: Regulatory frameworks such as the NAIC Model Bulletin on AI (2023) evaluate how insurers use AI internally—in claims processing, underwriting automation, and pricing models. AIRS evaluates the AI security risk of the entities those insurers choose to underwrite. The NAIC tool looks inward. Underwriting needs to look outward. AIRS bridges that gap.

In scope:

- Organizations deploying AI systems in production environments
- Organizations integrating third-party AI models, APIs, or services
- Organizations developing proprietary AI/ML models for internal or external use
- AI risk assessment for insurance underwriting, renewal, and portfolio management
- Regulatory compliance demonstration for AI-specific requirements

Out of scope:

- Assessment of insurer-internal AI governance (addressed by NAIC and similar frameworks)
- General IT security posture (addressed by existing cybersecurity rating methodologies)
- AI model performance benchmarking (accuracy, latency, throughput metrics)
- AI ethics evaluation beyond documented compliance and fairness testing requirements

1.4 Relationship to Other Standards and Frameworks

AIRS is designed to complement, not replace, existing regulatory frameworks and standards. The following table summarizes the relationship between AIRS and key frameworks referenced throughout this specification.

Table 1: Relationship to Referenced Standards

Framework	Relationship to AIRS	Reference
NIST AI RMF 1.0	AIRS domain structure aligns with AI RMF functions (Govern, Map, Measure, Manage). Cross-mapping provided in Annex B.	NIST AI 100-1 (2023)
ISO/IEC 42001:2023	AIRS assessment process aligns with ISO 42001 management system clauses. Factor-level crosswalk in Annex B.	ISO/IEC 42001:2023
EU AI Act	AIRS factors address EU AI Act high-risk system requirements. Regulatory alignment table in Annex B.	Regulation (EU) 2024/1689
CVSS v4.0	AIRS scoring architecture informed by CVSS metric presentation and severity scale conventions.	FIRST.Org (2023)

Framework	Relationship to AIRS	Reference
FAIR v3.0	AIRS risk quantification approach informed by FAIR ontological decomposition methodology.	FAIR Institute (2025)
NIST SP 800-53	AIRS cybersecurity factors reference SP 800-53 control families where applicable.	NIST SP 800-53r5 (2020)

1.5 Intended Audience

Table 2: Intended Audience and Use Cases

Audience	Primary Use of This Specification
Insurance Carriers	Implementing AIRS-based underwriting criteria, policy language, and risk selection
Reinsurers	Incorporating AIRS tiers into treaty terms, portfolio analysis, and accumulation monitoring
Regulators	Referencing AIRS methodology in bulletins, guidance, and examination procedures
Enterprise Risk Teams	Conducting self-assessments, identifying remediation priorities, demonstrating AI security maturity
Auditors and Assessors	Performing AIRS-conformant assessments using the rubrics and evidence standards defined herein
Academic Researchers	Referencing and extending the AIRS methodology in published research

1.6 Document Organization

This specification follows ISO-style clause numbering conventions. Sections 1–3 provide introductory, normative reference, and definitional material. Sections 4–10 contain the normative specification body. Annexes A through D provide supplementary material, with Annex A designated as normative (binding) and Annexes B–D as informative.

1.7 Notation Conventions

Table 3: Notation Conventions

Term	Meaning
shall	Indicates a mandatory requirement. Non-compliance with a 'shall' statement constitutes non-conformance with this standard.
should	Indicates a recommendation. Deviation from a 'should' statement is permissible with documented justification.
may	Indicates a permission or option. Implementation is at the discretion of the adopting organization.

Term	Meaning
normative	Content that is binding and subject to conformance requirements.
informative	Content that provides guidance, context, or explanation but is not subject to conformance requirements.

Section 2: Normative References

The following standards and frameworks are referenced normatively within this specification. Where a dated reference is cited, only the edition cited applies. Where an undated reference is cited, the latest edition of the referenced document applies.

[1] NIST AI 100-1 (2023). Artificial Intelligence Risk Management Framework (AI RMF 1.0). National Institute of Standards and Technology. https://doi.org/10.6028/NIST.AI.100-1

[2] ISO/IEC 42001:2023. Information technology — Artificial intelligence — Artificial intelligence management system. International Organization for Standardization.

[3] Regulation (EU) 2024/1689 of the European Parliament and of the Council. Laying down harmonised rules on artificial intelligence (Artificial Intelligence Act). Official Journal of the European Union.

[4] NIST SP 800-53, Revision 5 (2020). Security and Privacy Controls for Information Systems and Organizations. National Institute of Standards and Technology. https://doi.org/10.6028/NIST.SP.800-53r5

[5] FIRST.Org, Inc. (2023). Common Vulnerability Scoring System version 4.0: Specification Document. https://www.first.org/cvss/v4.0/specification-document

[6] FAIR Institute (2025). Factor Analysis of Information Risk (FAIR) Standard, Version 3.0. FAIR Institute. https://www.fairinstitute.org/

[7] ISO/IEC 22989:2022. Information technology — Artificial intelligence — Artificial intelligence concepts and terminology. International Organization for Standardization.

Section 3: Terms and Definitions

For the purposes of this specification, the following terms and definitions apply. Where a term is defined in ISO/IEC 22989:2022 or NIST AI 100-1, the source is noted.

adversarial robustness
The degree to which an AI system maintains correct function when subjected to intentionally crafted inputs designed to cause misclassification, evasion, or degradation.

AI Insurance Readiness Score (AIRS)
A composite risk score on a 0–100 scale derived from weighted assessment of five scoring domains, quantifying an entity's AI security maturity for insurance underwriting purposes.

AI supply chain
The set of third-party providers, foundation models, APIs, data sources, infrastructure services, and sub-processors upon which an organization's AI systems depend.

assessment period
The time window during which evidence is collected and evaluated for an AIRS assessment. The standard assessment period is twelve months.

assessed entity
The organization whose AI security posture is being evaluated under this standard.

assessor
An individual or organization performing an AIRS assessment in accordance with the requirements of Section 8.

blast radius
The scope of impact resulting from an AI system failure, measured in terms of affected users, data, systems, or business processes.

composite score
The final AIRS score calculated by applying domain weights to individual domain averages, expressed on a 0–100 scale. See Section 7.

concentration risk
The risk arising from dependency on a single or small number of AI providers, models, or infrastructure components.

conformant assessment
An AIRS assessment conducted in full compliance with the mandatory requirements specified in Section 9.

data provenance
The documented chain of custody, origin, licensing, and consent status of data used in AI model training, fine-tuning, or inference.

domain
One of five thematic groupings of risk factors defined in Section 5. Each domain receives a weight in the composite score calculation.

domain average
The arithmetic mean of the five factor scores within a single domain, expressed on a 1–5 scale.

drift
A change in the statistical properties of model inputs or outputs over time that may degrade model performance. Includes concept drift, data drift, and prediction drift.

factor
A discrete, measurable aspect of AI security risk within a domain. AIRS defines twenty-five factors, five per domain.

factor score
A maturity rating assigned to a single factor on a 1–5 scale, based on the rubric definitions in Annex A.

foundation model
A large-scale AI model trained on broad data that can be adapted to a wide range of downstream tasks. Includes large language models (LLMs), multimodal models, and similar architectures.

human-in-the-loop
A system design pattern in which human judgment is required at defined decision points before AI outputs are acted upon in high-stakes contexts.

maturity level
One of five ordinal levels (1–5) characterizing the sophistication and consistency of an organization's controls for a given factor. See Section 5, Table 6.

model quarantine
The process of isolating an AI model from production use pending investigation of a suspected failure, vulnerability, or adversarial compromise.

normative
Content within this specification that is binding and subject to conformance requirements. Deviations from normative content constitute non-conformance.

output liability
Legal and financial exposure arising from AI-generated content, decisions, or recommendations that cause harm to third parties.

rating tier
One of four ordinal classifications (AI Insurance Ready, Conditionally Insurable, Elevated Risk, Uninsurable) assigned based on composite AIRS score ranges. See Section 6.

reassessment trigger
A material change in the assessed entity's AI operations, risk profile, or regulatory environment that necessitates a new AIRS assessment before the standard reassessment period.

risk
The probable frequency and probable magnitude of future loss arising from AI system deployment, operation, or failure. Aligned with FAIR v3.0 definition.

scoring rubric
The set of maturity level definitions and evaluative criteria for a specific factor, as specified in Annex A.

sub-processor
A third party that processes data on behalf of an AI vendor or the assessed entity as part of the AI pipeline.

systemic resilience
The capacity of an organization's AI infrastructure to withstand, adapt to, and recover from failures, attacks, or disruptions without cascading impact.

underwriting implication
The recommended insurance coverage decision, pricing adjustment, or policy term modification associated with a given AIRS rating tier.

weight
A decimal multiplier applied to a domain average in the composite score formula, reflecting the domain's relative importance to overall AI security risk.

Section 4: Framework Overview

4.1 Framework Design Rationale

AIRS was designed according to five guiding principles that distinguish it from existing cybersecurity rating methodologies:

• AI-native risk coverage. Every factor in the AIRS framework addresses a risk specific to or materially amplified by artificial intelligence. Traditional IT hygiene metrics are excluded from the scoring model, though they may inform contextual assessment.

• Transparency. The complete scoring methodology—domains, factors, weights, rubrics, and tier definitions—is published as an open standard. No proprietary black-box scoring. Any qualified assessor can reproduce the scoring methodology independently.

• Insurance alignment. AIRS is calibrated to the decision architecture of institutional underwriting. Rating tiers map directly to coverage decisions, premium adjustments, and policy terms. The framework speaks the language of carriers and reinsurers, not only technologists.

• Regulatory compatibility. Each AIRS factor maps to specific provisions of NIST AI RMF 1.0, ISO/IEC 42001:2023, and the EU AI Act. Organizations already investing in regulatory compliance can leverage existing evidence for AIRS assessment.

• Maturity orientation. AIRS measures the maturity and consistency of controls, not their mere existence. A five-level maturity scale distinguishes between organizations with documented policies and those with continuously monitored, optimized practices.

4.2 Scoring Architecture

The AIRS scoring architecture comprises three layers:

Layer 1: Factor Assessment. Twenty-five individual risk factors are each assessed on a 1–5 maturity scale using the rubric definitions specified in Annex A. Each factor score reflects the assessed entity's current maturity for that specific risk dimension.

Layer 2: Domain Aggregation. Factor scores are grouped into five domains. Each domain average is the arithmetic mean of its five constituent factor scores, expressed on a 1–5 scale.

Layer 3: Composite Scoring. Domain averages are combined through a weighted sum, then scaled by a factor of 20 to produce the final AIRS composite score on a 0–100 scale. The composite score maps to one of four rating tiers.

Table 4: AIRS Scoring Architecture Summary

Component	Count	Scale	Method
Risk Factors	25 (5 per domain)	1–5 maturity	Rubric-based assessment

Component	Count	Scale	Method
Scoring Domains	5	1.0–5.0 average	Arithmetic mean of 5 factors
Domain Weights	5	0.15–0.25	Fixed (see Section 7)
Composite Score	1	0–100	Weighted sum × 20
Rating Tiers	4	Ordinal	Score range mapping

4.3 Rating Scale Overview

AIRS produces a composite score on a 0–100 scale that maps to four rating tiers. The following table provides an overview; detailed tier definitions appear in Section 6.

Table 5: AIRS Rating Tier Overview

AIRS Score Range	Rating Tier	Classification	Underwriting Signal
80–100	Tier 1	AI Insurance Ready	Broadest coverage, premium discounts
60–79	Tier 2	Conditionally Insurable	Coverage with exclusions/sub-limits
40–59	Tier 3	Elevated Risk	Limited coverage, mandatory remediation
20–39	Tier 4	Uninsurable	Coverage denied pending remediation

NOTE: The theoretical minimum AIRS score is 20 (all factors scored at maturity level 1). The formula produces a minimum of 20, not 0, because even the lowest maturity ratings contribute a base score. See Section 7 for the complete formula.

4.4 Framework Governance

4.4.1 Standard Stewardship

AI Security Intelligence LLC serves as the steward of the AIRS standard. The steward is responsible for maintaining the specification, processing errata, managing the version lifecycle, and convening advisory review when standard revisions are considered.

4.4.2 Review and Update Cycle

The AIRS standard shall be reviewed at minimum every twenty-four months to assess whether updates are warranted based on changes in the AI risk landscape, regulatory developments, or feedback from institutional adopters. Minor updates (document version increments) may be published as needed for editorial corrections. Major updates (standard version increments) shall include a public comment period of not less than sixty days.

4.4.3 Errata and Feedback Process

Errata reports, implementation questions, and revision proposals may be submitted to standards@aisecurityintelligence.com. Confirmed errata shall be published in the version history table

and documented in Annex D.

4.5 Limitations and Appropriate Use

AIRS is one input to the underwriting decision process, not a substitute for actuarial judgment. The following limitations apply:

- AIRS scores reflect maturity at the time of assessment and are not predictive guarantees of future performance or incident avoidance.
- AIRS does not measure traditional IT security hygiene, which should be evaluated through complementary methodologies.
- AIRS factor weights reflect the standard's assessment of relative risk importance as of the publication date and may be revised in future versions.
- Assessment quality depends on the completeness and accuracy of evidence provided by the assessed entity and the competence of the assessor.
- AIRS scores are not directly comparable across different standard versions (e.g., v1.0 scores should not be compared to future v2.0 scores without normalization guidance).

Section 5: Scoring Domains

Domain Structure and Presentation Format

AIRS defines five scoring domains, each containing five risk factors. This section specifies each domain in a parallel structure: definition and rationale, constituent factors with descriptions, weight assignment, assessment guidance, and relationship to other domains.

Each factor is assessed on a 1–5 maturity scale. The general maturity level definitions that apply across all factors are specified in the following table. Factor-specific rubric detail is provided in Annex A (normative).

Table 6: General Maturity Level Definitions

Level	Rating	General Definition
1	Non-Existent	No documented controls, processes, or organizational awareness. The entity has not addressed this factor in any formal capacity.
2	Ad Hoc	Informal or inconsistent practices exist. Controls are reactive, undocumented, and dependent on individual knowledge rather than institutional process.
3	Defined	Documented policies and procedures exist. Controls are implemented but may not be consistently enforced, measured, or subjected to regular review.
4	Managed	Controls are consistently applied, measured against defined metrics, and reviewed on a regular cycle. Testing and improvement processes are established and functioning.
5	Optimized	Industry-leading practices with continuous improvement, automated monitoring, proactive threat anticipation, and demonstrated resilience under stress conditions.

5.1 Domain 1: Model Integrity and Validation

Weight: 25% (0.25)

5.1.1 Definition and Rationale

Model integrity is the foundation of AI security. An AI system whose models cannot be trusted—because training data is unverified, adversarial attacks succeed, version control is absent, poisoning goes undetected, or drift degrades performance—represents a fundamental risk that propagates through all downstream uses. This domain receives the highest weight (25%) because model-level failures can invalidate every other control in the risk stack.

5.1.2 Constituent Factors

Table 7: Domain 1 Factors

Factor	Name	Description
1.1	Training Data Provenance	Documentation of training data sources, lineage, licensing, and consent verification for all data used in model development.
1.2	Adversarial Robustness Testing	Frequency and rigor of adversarial testing protocols, including red-teaming cadence and vulnerability remediation timelines.
1.3	Model Versioning and Rollback	Capabilities for maintaining model version history, rapid rollback to known-good states, and change management procedures.
1.4	Poisoning Detection	Mechanisms for detecting data poisoning attacks, anomalous training inputs, and compromised data pipelines.
1.5	Drift Detection and Monitoring	Continuous monitoring for model performance degradation, concept drift, and distribution shift in production environments.

5.1.3 Factor Specifications

Factor 1.1: Training Data Provenance

> "Documentation of training data sources, lineage, licensing, and consent verification for all data used in model development."

Evaluates whether the assessed entity maintains a comprehensive data inventory for AI training datasets, including source documentation, licensing terms, consent chains, and data quality assessments. Higher maturity levels require automated provenance tracking, regular audits of data licensing compliance, and documented remediation processes for identified provenance gaps.

Factor 1.2: Adversarial Robustness Testing

> "Frequency and rigor of adversarial testing protocols, including red-teaming cadence and vulnerability remediation timelines."

Assesses the entity's program for testing AI models against adversarial inputs, evasion techniques, and attack scenarios. Evaluates the existence and cadence of red-team exercises, the breadth of attack vectors tested, remediation SLAs for identified vulnerabilities, and integration of adversarial testing into the model development lifecycle.

Factor 1.3: Model Versioning and Rollback

> "Capabilities for maintaining model version history, rapid rollback to known-good states, and change management procedures."

Evaluates the entity's infrastructure for tracking model versions, maintaining rollback capabilities, and governing model promotion decisions. Higher maturity levels require automated versioning, tested rollback procedures with documented recovery time objectives, and formal change management boards for production model updates.

Factor 1.4: Poisoning Detection

> "Mechanisms for detecting data poisoning attacks, anomalous training inputs, and compromised data pipelines."

Assesses controls for identifying and mitigating data poisoning—the deliberate introduction of malicious data into training pipelines. Evaluates anomaly detection on training inputs, integrity verification of data pipelines, monitoring for distribution anomalies, and incident response procedures specific to poisoning events.

Factor 1.5: Drift Detection and Monitoring

> "Continuous monitoring for model performance degradation, concept drift, and distribution shift in production environments."

Evaluates the entity's capabilities for detecting and responding to model drift—changes in input distributions, concept definitions, or model performance over time. Assesses monitoring infrastructure, alerting thresholds, automated retraining triggers, and documented procedures for drift remediation.

5.1.4 Assessment Guidance

Evidence for Domain 1 assessment typically includes: model development documentation, training data inventories, adversarial testing reports, version control system logs, monitoring dashboards and alert configurations, incident reports related to model failures, and change management records. Assessors should verify that documented policies are reflected in operational practice, not merely aspirational.

5.1.5 Relationship to Other Domains

Domain 1 has strong interdependencies with Domain 3 (AI Supply Chain Security), as supply chain vulnerabilities can compromise model integrity through poisoned training data or compromised model components. Domain 5 (Systemic Resilience) depends on Domain 1 controls for failure isolation and recovery capabilities.

5.2 Domain 2: Output Liability and Content Risk

Weight: 20% (0.20)

5.2.1 Definition and Rationale

AI systems generate outputs—text, decisions, recommendations, classifications—that create direct legal and financial liability for the deploying organization. Hallucinated content, biased decisions, and unvetted recommendations expose entities to litigation, regulatory action, and reputational harm. This domain evaluates the controls governing what AI systems produce and how organizations manage the liability those outputs create.

5.2.2 Constituent Factors

Table 8: Domain 2 Factors

Factor	Name	Description
2.1	Human-in-the-Loop Oversight	Protocols for human review and approval of high-stakes AI outputs, escalation triggers, and override mechanisms.
2.2	Output Validation	Automated and manual mechanisms for fact-checking, accuracy verification, and hallucination detection in AI-generated content.
2.3	Disclaimer and Disclosure	Frameworks for transparent disclosure of AI involvement, output limitations, and appropriate use guidance.
2.4	Error Correction and Recall	Procedures for correcting erroneous AI outputs, notifying affected parties, and documenting remediation actions.
2.5	Liability Documentation	Contractual liability allocation in terms of service, customer agreements, and vendor contracts governing AI outputs.

5.2.3 Factor Specifications

Factor 2.1: Human-in-the-Loop Oversight

> "Protocols for human review and approval of high-stakes AI outputs, escalation triggers, and override mechanisms."

Evaluates the extent to which the assessed entity maintains human oversight of AI-generated outputs in high-stakes contexts. Assesses the definition of 'high-stakes' thresholds, the design and enforcement of escalation triggers, override mechanisms, and documentation of human review decisions. Higher maturity levels require risk-calibrated oversight tiers with automated escalation routing.

Factor 2.2: Output Validation

> "Automated and manual mechanisms for fact-checking, accuracy verification, and hallucination detection in AI-generated content."

Assesses the entity's mechanisms for verifying the accuracy and reliability of AI outputs before they reach end users or inform business decisions. Evaluates hallucination detection systems, fact-checking workflows, confidence scoring, and quality assurance processes. Higher maturity levels require automated validation pipelines with defined accuracy thresholds and continuous monitoring.

Factor 2.3: Disclaimer and Disclosure

> "Frameworks for transparent disclosure of AI involvement, output limitations, and appropriate use guidance."

Evaluates the entity's transparency practices regarding AI use. Assesses whether AI involvement is disclosed to affected parties, whether output limitations are communicated, and whether appropriate use guidance accompanies AI-generated content. Higher maturity levels require context-specific disclosure frameworks calibrated to audience and risk level.

Factor 2.4: Error Correction and Recall

> "Procedures for correcting erroneous AI outputs, notifying affected parties, and documenting remediation actions."

Assesses the entity's ability to identify, correct, and communicate about erroneous AI outputs after they have been distributed. Evaluates recall procedures, notification protocols for affected parties, root cause analysis processes, and documentation of remediation actions. Higher maturity levels require automated error detection with defined response time objectives.

Factor 2.5: Liability Documentation

> "Contractual liability allocation in terms of service, customer agreements, and vendor contracts governing AI outputs."

Evaluates the entity's contractual framework for allocating liability arising from AI outputs. Assesses terms of service, customer agreements, vendor contracts, and indemnification provisions. Higher maturity levels require legal review cadence, jurisdiction-specific compliance, and documented liability allocation for all AI-enabled products and services.

5.2.4 Assessment Guidance

Evidence for Domain 2 assessment typically includes: human oversight policies and escalation logs, output validation system documentation and accuracy metrics, disclosure frameworks and user-facing notices, error correction incident reports, terms of service, and liability allocation provisions in customer and vendor agreements.

5.2.5 Relationship to Other Domains

Domain 2 is closely linked to Domain 4 (Regulatory and Compliance Alignment), as output liability controls directly support compliance with EU AI Act transparency requirements and U.S. state AI disclosure laws. Domain 1 (Model Integrity) provides the upstream controls that reduce output errors.

5.3 Domain 3: AI Supply Chain Security

Weight: 20% (0.20)

5.3.1 Definition and Rationale

Modern AI deployments depend on complex supply chains of foundation model providers, API services, data enrichment partners, and infrastructure vendors. A vulnerability or failure anywhere in this chain can propagate to the assessed entity. This domain evaluates the entity's visibility into, governance over, and resilience against supply chain risks specific to AI components.

5.3.2 Constituent Factors

Table 9: Domain 3 Factors

Factor	Name	Description
3.1	Vendor Due Diligence	Thoroughness of security, compliance, and risk assessment for foundation model providers and AI vendors.
3.2	Dependency Mapping	Comprehensive inventory of third-party models, APIs, and AI components with dependency chain documentation.
3.3	Sub-processor Auditing	Auditing protocols for data enrichment partners, sub-processors, and downstream data handlers in the AI pipeline.
3.4	Contractual Liability	Contractual allocation of liability with AI vendors including indemnification, SLAs, and breach notification requirements.
3.5	Data Provenance Verification	Mechanisms for verifying training data provenance, licensing compliance, and consent chains across the supply chain.

5.3.3 Factor Specifications

Factor 3.1: Vendor Due Diligence

> "Thoroughness of security, compliance, and risk assessment for foundation model providers and AI vendors."

Evaluates the entity's process for assessing AI vendors prior to engagement and on an ongoing basis. Assesses vendor security evaluation criteria, compliance verification, financial stability review, and documented approval processes. Higher maturity levels require continuous vendor monitoring, risk-rated vendor tiers, and contractual security requirements.

Factor 3.2: Dependency Mapping

> "Comprehensive inventory of third-party models, APIs, and AI components with dependency chain documentation."

Assesses the entity's visibility into its AI dependency chain. Evaluates whether all third-party models, APIs, libraries, and AI services are inventoried with documented dependency relationships. Higher maturity levels require automated dependency discovery, real-time inventory updates, and impact analysis capabilities for supply chain changes.

Factor 3.3: Sub-processor Auditing

> "Auditing protocols for data enrichment partners, sub-processors, and downstream data handlers in the AI pipeline."

Evaluates the entity's oversight of third parties that process data within the AI pipeline. Assesses audit rights, audit frequency, sub-processor inventory completeness, and documented audit findings with remediation tracking. Higher maturity levels require contractual audit rights, annual audit execution, and continuous compliance monitoring.

Factor 3.4: Contractual Liability

> "Contractual allocation of liability with AI vendors including indemnification, SLAs, and breach notification requirements."

Assesses the entity's contractual protections against AI vendor failures. Evaluates indemnification provisions, service level agreements with defined remedies, breach notification requirements, and data processing agreements. Higher maturity levels require standardized AI vendor contract templates with security-specific provisions reviewed by legal counsel.

Factor 3.5: Data Provenance Verification

> "Mechanisms for verifying training data provenance, licensing compliance, and consent chains across the supply chain."

Evaluates the entity's ability to verify that data used throughout the AI supply chain meets provenance, licensing, and consent requirements. Assesses upstream data verification procedures, licensing compliance tracking, and consent chain documentation for third-party data sources. Higher maturity levels require automated provenance verification and contractual provenance warranties from vendors.

5.3.4 Assessment Guidance

Evidence for Domain 3 assessment typically includes: vendor assessment documentation, approved vendor lists with risk ratings, dependency inventories, sub-processor registers, audit reports and findings, vendor contracts with AI-specific provisions, and data provenance verification records.

5.3.5 Relationship to Other Domains

Domain 3 directly supports Domain 1 (Model Integrity) by securing the inputs and components upon which model integrity depends. Domain 5 (Systemic Resilience) addresses the concentration and continuity risks that emerge from supply chain dependencies identified in Domain 3.

5.4 Domain 4: Regulatory and Compliance Alignment

Weight: 15% (0.15)

5.4.1 Definition and Rationale

The regulatory landscape for artificial intelligence is evolving rapidly across jurisdictions. The EU AI Act, U.S. state-level AI regulations, cross-border data transfer requirements, and algorithmic fairness mandates create compliance obligations that directly affect insurable risk. Non-compliance exposes entities to regulatory fines, litigation, and reputational harm—all of which represent quantifiable loss exposures for underwriters. This domain receives the lowest weight (15%) because regulatory requirements are a lagging indicator, codifying risks that the other four domains measure proactively.

5.4.2 Constituent Factors

Table 10: Domain 4 Factors

Factor	Name	Description
4.1	EU AI Act Readiness	Assessment of preparedness for EU AI Act high-risk requirements including conformity assessment, documentation, and human oversight.
4.2	U.S. State Law Compliance	Mapping of compliance posture across relevant U.S. state AI regulations, including California, Texas, and Colorado.
4.3	Cross-Border Data Transfer	Protocols for lawful cross-border data transfers including GDPR adequacy, standard contractual clauses, and data localization.
4.4	Bias Testing and Fairness	Documentation and cadence of algorithmic bias testing, fairness metrics, and disparate impact analysis.
4.5	Regulatory Reporting	Procedures for regulatory incident notification, mandatory reporting compliance, and records retention.

5.4.3 Factor Specifications

Factor 4.1: EU AI Act Readiness

> "Assessment of preparedness for EU AI Act high-risk requirements including conformity assessment, documentation, and human oversight."

Evaluates the entity's preparedness for EU AI Act obligations, focusing on high-risk AI system requirements. Assesses conformity assessment readiness, technical documentation completeness, human oversight mechanisms, and quality management system alignment. Higher maturity levels require proactive classification of AI systems against EU AI Act risk categories with documented compliance roadmaps.

Factor 4.2: U.S. State Law Compliance

> "Mapping of compliance posture across relevant U.S. state AI regulations, including California, Texas, and Colorado."

Assesses the entity's compliance mapping across applicable U.S. state AI regulations. Evaluates awareness of jurisdictional requirements, compliance gap analysis, and documented remediation plans. Higher maturity levels require continuous regulatory monitoring, jurisdiction-specific compliance programs, and legal counsel review of AI deployments.

Factor 4.3: Cross-Border Data Transfer

> "Protocols for lawful cross-border data transfers including GDPR adequacy, standard contractual clauses, and data localization."

Evaluates the entity's protocols for lawful cross-border transfer of AI-related data. Assesses GDPR adequacy mechanisms, standard contractual clauses, data localization compliance, and transfer impact assessments. Higher maturity levels require automated transfer compliance monitoring and jurisdiction-specific data routing.

Factor 4.4: Bias Testing and Fairness

> "Documentation and cadence of algorithmic bias testing, fairness metrics, and disparate impact analysis."

Assesses the entity's program for testing AI systems for algorithmic bias and ensuring fairness. Evaluates bias testing methodology, fairness metric selection, testing frequency, remediation processes, and documentation of disparate impact analysis. Higher maturity levels require continuous bias monitoring, multiple fairness metrics, and independent fairness audits.

Factor 4.5: Regulatory Reporting

> "Procedures for regulatory incident notification, mandatory reporting compliance, and records retention."

Evaluates the entity's procedures for meeting regulatory reporting obligations related to AI incidents. Assesses notification timelines, reporting completeness, records retention policies, and coordination with legal counsel. Higher maturity levels require automated regulatory reporting workflows with jurisdiction-specific templates and pre-approved notification language.

5.4.4 Assessment Guidance

Evidence for Domain 4 assessment typically includes: regulatory compliance matrices, EU AI Act classification documentation, state law compliance assessments, cross-border data transfer agreements, bias testing reports with fairness metrics, regulatory notification procedures, and records retention schedules.

5.4.5 Relationship to Other Domains

Domain 4 interacts with all other domains. Model integrity controls (Domain 1) support conformity assessment requirements. Output liability controls (Domain 2) directly enable compliance with transparency and disclosure mandates. Supply chain governance (Domain 3) supports cross-border data transfer compliance.

5.5 Domain 5: Systemic Resilience and Continuity

Weight: 20% (0.20)

5.5.1 Definition and Rationale

AI systems are increasingly critical infrastructure. When they fail—whether through adversarial attack, vendor outage, model degradation, or cascading dependency failure—the blast radius can extend across entire business operations. This domain evaluates the entity's capacity to prevent single points of failure, isolate AI system failures, maintain business continuity during AI disruptions, and respond effectively to AI-specific incidents.

5.5.2 Constituent Factors

Table 11: Domain 5 Factors

Factor	Name	Description
5.1	Model Diversification	Strategy for reducing foundation model concentration including multi-vendor approaches and fallback architectures.
5.2	Failure Isolation	Architecture for containing AI failures including circuit breakers, graceful degradation, and blast radius limitation.
5.3	Business Continuity	Business continuity and disaster recovery planning specific to AI system failures, including manual fallback procedures.
5.4	Concentration Monitoring	Ongoing monitoring of shared infrastructure dependencies, single points of failure, and concentration risk metrics.
5.5	AI Incident Response	Incident response playbooks specific to AI failures including model quarantine, output recall, and stakeholder notification.

5.5.3 Factor Specifications

Factor 5.1: Model Diversification

> "Strategy for reducing foundation model concentration including multi-vendor approaches and fallback architectures."

Evaluates the entity's strategy for reducing dependency on any single AI model or provider. Assesses multi-vendor approaches, model redundancy architectures, tested failover procedures, and documented diversification rationale. Higher maturity levels require active multi-model deployment with automated failover and regular diversification reviews.

Factor 5.2: Failure Isolation

> "Architecture for containing AI failures including circuit breakers, graceful degradation, and blast radius limitation."

Assesses the entity's architectural controls for containing AI system failures. Evaluates circuit breaker implementations, graceful degradation patterns, blast radius limitation measures, and failure propagation analysis. Higher maturity levels require automated circuit breakers with defined failure thresholds, tested degradation modes, and documented blast radius boundaries.

Factor 5.3: Business Continuity

> "Business continuity and disaster recovery planning specific to AI system failures, including manual fallback procedures."

Evaluates the entity's business continuity planning for AI-specific disruption scenarios. Assesses AI-specific BCP documentation, manual fallback procedures, recovery time objectives, and tested recovery procedures. Higher maturity levels require regular AI-specific BCP exercises, documented recovery playbooks, and validated manual fallback capabilities.

Factor 5.4: Concentration Monitoring

> "Ongoing monitoring of shared infrastructure dependencies, single points of failure, and concentration risk metrics."

Assesses the entity's visibility into concentration risk across AI infrastructure. Evaluates monitoring of shared dependencies, single point of failure identification, concentration risk metrics, and documented mitigation strategies. Higher maturity levels require real-time concentration dashboards, automated alerting on threshold breaches, and regular concentration risk reviews.

Factor 5.5: AI Incident Response

> "Incident response playbooks specific to AI failures including model quarantine, output recall, and stakeholder notification."

Evaluates the entity's incident response capabilities for AI-specific incidents. Assesses playbook existence and completeness, model quarantine procedures, output recall capabilities, stakeholder notification protocols, and post-incident review processes. Higher maturity levels require tested IR playbooks, regular tabletop exercises, automated quarantine capabilities, and documented lessons-learned integration.

5.5.4 Assessment Guidance

Evidence for Domain 5 assessment typically includes: architecture documentation showing redundancy and isolation controls, business continuity plans with AI-specific scenarios, concentration risk dashboards, incident response playbooks, tabletop exercise reports, and post-incident review records.

5.5.5 Relationship to Other Domains

Domain 5 depends on supply chain visibility from Domain 3 for effective concentration monitoring and diversification planning. Domain 1 (Model Integrity) provides the versioning and rollback capabilities that enable rapid recovery from model-level incidents.

Section 6: Rating Tiers

6.1 Tier Design Principles

AIRS rating tiers translate composite scores into actionable underwriting classifications. The tier system is designed according to three principles:

- Actionability: Each tier maps to a distinct underwriting decision pathway—not a spectrum of possibilities, but a clear signal for coverage, pricing, and terms.
- Remediation orientation: Tiers are not permanent labels. The tier system is designed to incentivize improvement by providing a visible target for the assessed entity and a measurable baseline for the carrier.
- Institutional clarity: Tier names and definitions use language familiar to insurance professionals. The four-tier structure mirrors established insurance rating conventions without adopting letter grades or opaque numerical bands.

6.2 Tier Definitions

The following definitions are normative. Conformant AIRS implementations shall apply these tier definitions without modification.

TIER 1: AI INSURANCE READY

AIRS Score Range: 80–100

The assessed entity demonstrates comprehensive, mature AI security controls across all five domains. Controls are consistently applied, measured, reviewed, and improved. The entity maintains industry-leading practices in critical areas and demonstrates resilience under stress conditions.

Key Characteristics:

- Documented and enforced policies across all 25 factors with evidence of regular review and update
- Automated monitoring and alerting for model integrity, drift, and supply chain changes
- Tested incident response playbooks with demonstrated model quarantine and output recall capabilities
- Comprehensive regulatory compliance program with proactive jurisdictional monitoring
- Multi-vendor AI strategy with tested failover and recovery procedures
- Regular third-party or independent assessments validating control effectiveness

Underwriting Implication: Qualifies for affirmative AI coverage with the broadest available terms. Premium discounts of 10–25% relative to base rates. Eligible for highest coverage limits and fewest

exclusions. Recommended for preferred risk classification.

TIER 2: CONDITIONALLY INSURABLE

AIRS Score Range: 60–79

The assessed entity demonstrates established AI security controls with documented policies and implemented procedures across most domains. Controls are generally consistent but may lack automation, continuous monitoring, or demonstrated resilience. Specific domains may show maturity gaps that require targeted remediation.

Key Characteristics:

- Documented policies exist for most factors but may lack consistent enforcement or measurement
- Manual or periodic monitoring processes in place but automated continuous monitoring is limited
- Incident response plans documented but may not be regularly tested or exercised
- Regulatory compliance efforts underway but gaps may exist in emerging jurisdictions
- Some supply chain visibility but dependency mapping may be incomplete
- Remediation roadmap identifiable with achievable milestones to Tier 1

Underwriting Implication: Coverage available with specific exclusions, sub-limits, and higher premiums reflecting identified gaps. Remediation milestones may be incorporated as policy conditions, with premium adjustments available upon demonstrated improvement. Coverage terms may restrict AI-related endorsements pending remediation.

TIER 3: ELEVATED RISK

AIRS Score Range: 40–59

The assessed entity demonstrates basic awareness of AI security requirements but significant gaps exist in control implementation, consistency, and maturity. Controls may be ad hoc, reactive, or dependent on individual knowledge rather than institutional process.

Key Characteristics:

- Policies may exist in some areas but are inconsistently implemented or enforced
- Limited or no automated monitoring of AI systems in production
- Incident response for AI-specific scenarios is informal or untested
- Regulatory compliance awareness is emerging but systematic programs are absent
- Minimal supply chain visibility with limited vendor assessment processes
- Significant remediation effort required to achieve Tier 2 status

Underwriting Implication: Limited coverage only, with significant exclusions for AI-related claims, reduced limits, and premiums reflecting elevated exposure. Mandatory remediation requirements with defined milestones and timelines. Coverage renewal contingent on demonstrated progress toward Tier 2 minimum thresholds.

TIER 4: UNINSURABLE

AIRS Score Range: 20–39

The assessed entity demonstrates minimal or no AI security controls. AI systems are deployed without documented governance, monitoring, or risk management. The absence of fundamental controls creates risk exposures that exceed acceptable thresholds for insurance coverage.

Key Characteristics:

- No documented AI security policies or governance framework
- No monitoring of AI systems in production for integrity, drift, or adversarial activity
- No incident response capability specific to AI failures
- No awareness of or preparation for AI-specific regulatory requirements
- No visibility into AI supply chain dependencies or concentration risks
- Fundamental remediation required across all domains before coverage consideration

Underwriting Implication: Coverage denied until remediation milestones are achieved and a subsequent AIRS assessment demonstrates improvement to Tier 3 minimum. The assessed entity shall be provided with a remediation roadmap identifying priority actions to reach the Tier 3 threshold.

6.3 Tier Assignment Rules

6.3.1 Composite Score to Tier Mapping

Tier assignment is determined by the composite AIRS score as specified in Table 12. Scores falling on a tier boundary shall be assigned to the higher tier (e.g., a score of exactly 80 is assigned to Tier 1).

Table 12: Composite Score to Tier Mapping

Composite AIRS Score	Assigned Tier	Classification
80.0 – 100.0	Tier 1	AI Insurance Ready
60.0 – 79.9	Tier 2	Conditionally Insurable
40.0 – 59.9	Tier 3	Elevated Risk
20.0 – 39.9	Tier 4	Uninsurable

6.3.2 Domain Floor Rules

In addition to the composite score threshold, the following domain floor rules apply. An entity may not be assigned a tier higher than warranted by its weakest domain:

- Tier 1 eligibility: No single domain average below 3.5.
- Tier 2 eligibility: No single domain average below 2.5.

If a domain floor rule prevents assignment to a tier that the composite score would otherwise warrant, the entity shall be assigned to the highest tier for which all conditions are met. The limiting domain shall be identified in the assessment report.

6.3.3 Override and Exceptional Circumstances

This standard does not define override mechanisms for tier assignment. The composite score and domain floor rules determine the tier. Carriers may apply additional underwriting judgment to coverage decisions within a tier, but shall not represent a non-conformant tier assignment as AIRS-conformant.

6.4 Tier Progression and Reassessment

Tier progression is achieved through remediation of identified control gaps and subsequent reassessment. The AIRS framework encourages improvement through three mechanisms:

- Remediation roadmaps: Each AIRS assessment should produce a prioritized list of remediation actions that would most efficiently improve the composite score and tier classification.
- Interim assessments: Assessed entities may request reassessment at any time to demonstrate improvement. The standard assessment period of twelve months does not preclude earlier reassessment.
- Conditional milestones: Carriers may incorporate specific remediation milestones as policy conditions, enabling premium adjustments or expanded coverage terms upon demonstrated completion.

Section 7: Composite Score Calculation

7.1 Composite Score Formula

The AIRS composite score is calculated as the weighted sum of domain averages, scaled to a 0–100 range. The formula is normative; conformant implementations shall apply it without modification.

AIRS Composite Score Formula:

$$AIRS = (D_1 \times w_1 + D_2 \times w_2 + D_3 \times w_3 + D_4 \times w_4 + D_5 \times w_5) \times 20$$

Where:

D_i = arithmetic mean of the five factor scores in Domain i (range: 1.0 – 5.0)

w_i = weight assigned to Domain i (see Table 13)

The scaling factor of 20 maps the weighted domain average (range: 1.0 – 5.0) to a 0 – 100 scale.

Table 13: Domain Weights

Domain	Weight (w)	Percentage	Rationale
D1: Model Integrity and Validation	0.25	25%	Foundational risk; model failures propagate through all downstream controls
D2: Output Liability and Content Risk	0.20	20%	Direct legal/financial exposure from AI-generated outputs
D3: AI Supply Chain Security	0.20	20%	Third-party dependency risks specific to AI component chains
D4: Regulatory and Compliance Alignment	0.15	15%	Lagging indicator; codifies risks measured proactively in other domains
D5: Systemic Resilience and Continuity	0.20	20%	Infrastructure-level failure and concentration risks

NOTE: Domain weights sum to 1.00. Conformant implementations shall not modify these weights. The standard reserves weight adjustment for future versions based on empirical claims data and evolving risk landscapes.

7.2 Normalization and Scale

Each domain average D_i is calculated as the arithmetic mean of the five constituent factor scores within that domain:

$$D_i = (F_{i,1} + F_{i,2} + F_{i,3} + F_{i,4} + F_{i,5}) / 5$$

Where $F_{i,j}$ is the factor score for Factor j in Domain i (integer, range 1–5).

The composite score range is determined by the factor score boundaries:

Table 14: Score Range Boundaries

Scenario	All Factor Scores	Domain Averages	Weighted Sum	AIRS Score
Minimum	1 (all factors)	1.0 (all domains)	1.0	20
Maximum	5 (all factors)	5.0 (all domains)	5.0	100
Midpoint	3 (all factors)	3.0 (all domains)	3.0	60

7.3 Worked Examples

Example 1: Tier 1 Organization

A technology company with mature AI governance receives the following domain averages:

Domain	Factor Scores	Domain Average (D)	Weight (w)	Weighted Contribution
D1: Model Integrity	5, 4, 5, 4, 4	4.40	0.25	1.100
D2: Output Liability	4, 4, 5, 4, 4	4.20	0.20	0.840
D3: Supply Chain	5, 4, 4, 4, 5	4.40	0.20	0.880
D4: Regulatory	4, 3, 4, 4, 4	3.80	0.15	0.570
D5: Resilience	4, 5, 4, 4, 5	4.40	0.20	0.880
			Weighted Sum:	4.270
			AIRS Score:	85.4

AIRS = 4.270 × 20 = 85.4 — Tier 1: AI Insurance Ready. All domain averages exceed 3.5 floor rule. Tier 1 assignment confirmed.

Example 2: Domain Floor Rule Application

An entity with strong controls in most areas but weak supply chain governance:

Domain	Factor Scores	Domain Average (D)	Weight (w)	Weighted Contribution
D1: Model Integrity	5, 5, 4, 4, 5	4.60	0.25	1.150
D2: Output Liability	4, 5, 4, 4, 4	4.20	0.20	0.840
D3: Supply Chain	2, 2, 1, 2, 2	1.80	0.20	0.360
D4: Regulatory	4, 3, 4, 4, 4	3.80	0.15	0.570
D5: Resilience	5, 4, 4, 5, 4	4.40	0.20	0.880

Domain	Factor Scores	Domain Average (D)	Weight (w)	Weighted Contribution
			Weighted Sum:	3.800
			AIRS Score:	76.0

AIRS = 3.800 × 20 = 76.0. Composite score qualifies for Tier 2 (60–79). However, Domain 3 average (1.80) is below the Tier 2 floor of 2.5. Domain floor rule applies: entity is assigned Tier 3 (Elevated Risk), with Domain 3 identified as the limiting domain in the assessment report.

7.4 Score Precision and Rounding

AIRS composite scores shall be calculated to one decimal place and reported as such (e.g., 76.0, not 76). Domain averages shall be calculated to two decimal places for intermediate computation and may be reported at one or two decimal places. Factor scores are integers (1–5) and shall not be reported with decimal precision.

Section 8: Assessment Process

8.1 Assessment Framework

An AIRS assessment is a structured evaluation of an entity's AI security posture across all five scoring domains. This section defines the requirements for conducting assessments that conform to the AIRS standard. The assessment framework is designed to be implementable by qualified assessors using documented evidence, without dependency on any specific assessment tool or platform.

8.2 Roles and Responsibilities

8.2.1 Assessor Requirements

A conformant AIRS assessment shall be conducted by an assessor meeting the following minimum qualifications:

- Demonstrated knowledge of AI/ML system architectures, development lifecycles, and deployment patterns
- Familiarity with cybersecurity risk assessment methodologies (e.g., NIST CSF, ISO 27001)
- Understanding of insurance underwriting principles and risk classification
- Ability to evaluate documentary evidence against the maturity rubrics defined in Annex A
- Independence from the assessed entity for third-party assessments (self-assessments are permitted but shall be disclosed)

Assessor qualification criteria are designed to establish a professional standard analogous to those governing independent controls examinations, emphasizing subject matter expertise, independence, and ongoing professional competency.

8.2.2 Assessed Entity Requirements

The assessed entity shall:

- Designate a point of contact responsible for coordinating the assessment
- Provide access to documentation, systems, and personnel necessary for evidence evaluation
- Disclose all AI systems in scope for the assessment, including third-party and embedded AI
- Attest to the completeness and accuracy of evidence provided

8.3 Evidence Standards

8.3.1 Required Documentation

Factor scores shall be supported by documentary evidence demonstrating the assessed entity's current practices. The following evidence categories apply:

Table 15: Evidence Categories

Evidence Category	Examples	Applicable Domains
Governance documentation	AI security policies, governance charter, risk management framework, board reporting	All domains
Technical documentation	Architecture diagrams, model cards, API documentation, monitoring configurations	D1, D3, D5
Process documentation	Procedures, runbooks, escalation matrices, change management records	All domains
Testing and audit records	Red-team reports, bias testing results, vendor audit findings, BCP exercise reports	D1, D2, D4, D5
Contractual documentation	Vendor agreements, terms of service, data processing agreements, SLAs	D2, D3
Compliance records	Regulatory compliance matrices, certification evidence, regulatory correspondence	D4
Operational evidence	Monitoring dashboards, alert logs, incident reports, remediation records	D1, D2, D5

8.3.2 Evidence Quality Criteria

Evidence submitted for AIRS assessment shall meet the following quality criteria:

- Currency: Evidence shall be dated within the assessment period (twelve months preceding the assessment date) unless demonstrating a durable control (e.g., architecture design).
- Specificity: Evidence shall specifically address the factor being assessed, not merely reference a general capability.
- Verifiability: Evidence shall be independently verifiable through documentation review, system inspection, or attestation.
- Completeness: Evidence shall address the full scope of the factor definition, not only selected elements.

8.3.3 Handling Insufficient Evidence

Where the assessed entity cannot provide evidence for a factor, the assessor shall assign a maturity level of 1 (Non-Existent) for that factor. The assessment report shall document which factors were scored at Level 1 due to insufficient evidence versus verified absence of controls.

8.4 Assessment Stages

A conformant AIRS assessment shall include the following stages:

Stage 1: Scoping. Define the assessment boundary, identify AI systems in scope, confirm the assessment period, and establish evidence collection requirements.

Stage 2: Evidence Collection. Gather documentation, conduct stakeholder interviews, review systems and configurations, and compile the evidence portfolio for each factor.

Stage 3: Factor Scoring. Evaluate evidence against the maturity rubrics in Annex A and assign a maturity level (1–5) to each of the twenty-five factors.

Stage 4: Domain Aggregation. Calculate domain averages as the arithmetic mean of the five factor scores within each domain.

Stage 5: Composite Scoring. Apply the domain weights and scaling factor to produce the composite AIRS score per the formula in Section 7.

Stage 6: Tier Assignment. Map the composite score to a rating tier per the rules in Section 6, including domain floor rule evaluation.

Stage 7: Reporting. Produce the AIRS assessment report documenting scores, evidence summaries, tier assignment, domain floor rule application, and remediation recommendations.

8.5 Assessment Period and Currency

The standard AIRS assessment period is twelve months. An AIRS score is considered current for twelve months from the assessment completion date. After twelve months, the score is considered stale and shall not be represented as current without reassessment.

8.6 Reassessment Triggers

In addition to the standard twelve-month reassessment cycle, a reassessment should be triggered by any of the following material changes:

- Adoption of a new foundation model or significant change to existing model architecture
- Material change to AI supply chain (new primary vendor, vendor breach, vendor termination)
- AI-related security incident, data breach, or regulatory enforcement action
- Entry into a new jurisdiction with AI-specific regulatory requirements
- Significant change in AI deployment scope (new high-risk use case, new customer segment)
- Merger, acquisition, or organizational restructuring affecting AI governance

Section 9: Conformance

9.1 Conformance Scope

This section defines what constitutes conformant use of the AIRS standard. Conformance requirements apply to organizations conducting AIRS assessments, publishing AIRS scores, or referencing AIRS methodology in policy language, regulatory filings, or commercial communications.

9.2 Types of Conformance

9.2.1 Full Conformance

An AIRS assessment achieves full conformance when it satisfies all mandatory requirements specified in Section 9.3. Only fully conformant assessments may use the designation "AIRS v1.0 Conformant Assessment" in reporting and communications.

9.2.2 Partial Conformance

An assessment that applies the AIRS methodology with documented deviations from one or more mandatory requirements achieves partial conformance. Partially conformant assessments shall use the designation "Based on AIRS v1.0 methodology" and shall disclose all deviations in the assessment report.

9.2.3 Non-Conformant Reference

Organizations may reference AIRS concepts, domain definitions, or factor descriptions for informational purposes without conducting a conformant assessment. Such references shall not imply that an AIRS assessment has been performed and shall not include AIRS tier assignments.

9.3 Mandatory Requirements for Conformant Use

A conformant AIRS assessment shall satisfy all of the following requirements:

9.3.1 The assessment shall evaluate all twenty-five factors across all five scoring domains as defined in Section 5. No factor or domain may be omitted.

9.3.2 Factor scores shall be assigned using the five-level maturity scale defined in Section 5 (Table 6) and detailed in Annex A.

9.3.3 Domain averages shall be calculated as the arithmetic mean of the five constituent factor scores.

9.3.4 The composite AIRS score shall be calculated using the formula and weights specified in Section 7, without modification.

9.3.5 Rating tier assignment shall follow the rules specified in Section 6, including domain floor rules.

9.3.6 The assessment shall be conducted in accordance with the process requirements specified in Section 8.

9.3.7 The assessment report shall identify the standard version used (AIRS v1.0) and the document version of this specification referenced.

9.3.8 Any published AIRS score shall include the tier assignment, the composite score, and the assessment date.

9.3.9 The assessor shall document the evidence basis for each factor score.

9.3.10 Self-assessments shall be disclosed as such and shall not be represented as independent third-party assessments.

9.4 Permitted Adaptations

The following adaptations are permitted within a conformant assessment:

• Addition of supplementary factors or domains beyond the twenty-five specified factors, provided the core AIRS scoring is preserved and supplementary scores are reported separately.

• Use of proprietary evidence collection tools, assessment platforms, or automation, provided the scoring methodology remains unchanged.

• Development of industry-specific assessment guidance for particular sectors (e.g., healthcare AI, financial services AI), provided all twenty-five factors are assessed.

• Translation of this specification into other languages, provided the normative content is faithfully translated.

9.5 Prohibited Modifications

The following modifications are prohibited in a conformant assessment:

• Modification of domain weights from the values specified in Section 7.

• Omission of any factor or domain from the assessment.

• Alteration of tier score ranges or tier definitions from those specified in Section 6.

• Removal or modification of domain floor rules.

• Substitution of alternative maturity scales for the five-level scale defined in this standard.

• Publication of AIRS tier assignments without supporting composite scores.

• Use of the "AIRS v1.0 Conformant Assessment" designation for assessments with unresolved non-conformances.

Section 10: How to Reference This Standard

10.1 Official Standard Name and Identifier

Table 16: Official Standard Identifiers

Attribute	Value
Full Name	AI Insurance Readiness Score Open Standard Specification
Standard Identifier	AIRS v1.0
Document Version	1.0
Short Reference	AIRS v1.0
Publisher	AI Security Intelligence LLC
Persistent URL	https://www.aisecurityintelligence.com/airs-v1-specification

10.2 Citation Formats

APA Style:

> AI Security Intelligence LLC. (2026). AI Insurance Readiness Score Open Standard Specification (AIRS v1.0, Document Version 1.0). AI Security Intelligence LLC. https://www.aisecurityintelligence.com/airs-v1-specification

IEEE Style:

> AI Security Intelligence LLC, "AI Insurance Readiness Score Open Standard Specification," AIRS v1.0, Apr. 2026. [Online]. Available: https://www.aisecurityintelligence.com/airs-v1-specification

Policy Language Reference:

> "...as measured by the AI Insurance Readiness Score (AIRS) v1.0, published by AI Security Intelligence LLC (April 2026)."

Regulatory Filing Reference:

> "AI risk assessment conducted in accordance with the AI Insurance Readiness Score (AIRS) Open Standard Specification, Version 1.0 (AI Security Intelligence LLC, 2026)."

10.3 Attribution Requirements

Organizations publishing AIRS scores, referencing AIRS tiers in underwriting criteria, or citing the AIRS methodology in regulatory communications shall include the following elements in their attribution:

- The standard name: AI Insurance Readiness Score (AIRS)
- The standard version: v1.0
- The publisher: AI Security Intelligence LLC
- Whether the assessment is self-assessed or independently assessed
- The assessment date

10.4 Permitted Use Statement

The AIRS methodology is published as an open standard. The following uses are expressly permitted without additional licensing:

- Internal risk assessment using the AIRS methodology
- Integration of AIRS scoring into underwriting workflows and policy language
- Reference to AIRS methodology in regulatory filings and examination responses
- Academic research, citation, and extension of the AIRS framework
- Development of AIRS-conformant assessment tools and platforms
- Translation of this specification for use in other jurisdictions

Reproduction or redistribution of this specification document in its entirety requires prior written permission from AI Security Intelligence LLC. Excerpts and quotations for the purposes listed above are permitted with attribution.

Annex A: Scoring Rubrics (Normative)

A.1 Purpose and Structure

This annex provides the detailed scoring rubrics for all twenty-five AIRS factors. These rubrics are normative—conformant assessments shall use these definitions when assigning maturity levels. Each rubric table specifies the observable criteria that distinguish each maturity level for a given factor.

The rubrics follow a consistent five-level maturity progression:

- Level 1 (Non-Existent): No evidence of capability or awareness
- Level 2 (Ad Hoc): Informal, inconsistent, individual-dependent practices
- Level 3 (Defined): Documented policies and procedures, implemented but not measured
- Level 4 (Managed): Consistent application, measurement, regular review cycles
- Level 5 (Optimized): Continuous improvement, automation, proactive anticipation

A.2 Domain 1: Model Integrity and Validation

Table 17: Rubric for Factor 1.1 — Training Data Provenance

Level	Rating	Criteria
1	Non-Existent	No documentation of training data sources, lineage, or licensing. No awareness of data provenance as a risk factor.
2	Ad Hoc	Some informal records of data sources exist but are incomplete, inconsistently maintained, or limited to certain models. No systematic provenance tracking.
3	Defined	Documented data inventory exists for training datasets with source, licensing, and consent information. Inventory is maintained but may not be regularly audited or automated.
4	Managed	Comprehensive data inventory with regular audits of provenance accuracy. Automated tracking of data lineage across the ML pipeline. Remediation process for identified provenance gaps.
5	Optimized	Fully automated provenance tracking with real-time lineage documentation across all data pipelines. Regular third-party audits. Proactive monitoring for licensing changes and consent revocations. Industry-leading data governance framework.

Table 18: Rubric for Factor 1.2 — Adversarial Robustness Testing

Level	Rating	Criteria
1	Non-Existent	No adversarial testing of AI models. No awareness of adversarial attack vectors as a risk factor.

Level	Rating	Criteria
2	Ad Hoc	Some informal testing conducted reactively (e.g., after a reported issue). No regular cadence, no documented methodology, limited attack vector coverage.
3	Defined	Documented adversarial testing program with defined methodology and scheduled cadence (e.g., quarterly). Coverage of major attack vectors. Findings documented but remediation timelines may be informal.
4	Managed	Regular adversarial testing with comprehensive attack vector coverage. Defined remediation SLAs. Integration with model development lifecycle. Red-team exercises conducted at least annually. Metrics tracked.
5	Optimized	Continuous adversarial testing integrated into CI/CD pipeline. Automated testing across comprehensive attack taxonomy. Strict remediation SLAs with executive escalation. External red-team engagements. Proactive threat intelligence integration.

Table 19: Rubric for Factor 1.3 — Model Versioning and Rollback

Level	Rating	Criteria
1	Non-Existent	No model version control. No rollback capability. Models are deployed without version tracking.
2	Ad Hoc	Informal version tracking (e.g., file naming conventions). Rollback possible but untested and without documented procedure. No formal change management.
3	Defined	Model version control system in place with documented versioning policy. Rollback procedures documented. Change management process exists for production model updates.
4	Managed	Automated model versioning with tested rollback procedures. Documented recovery time objectives. Formal change management board for production deployments. Regular rollback testing.
5	Optimized	Fully automated version control with continuous deployment pipelines. Rollback tested as part of regular operations. Immutable model registries. Automated canary/blue-green deployment patterns. Change management integrated with risk assessment.

Table 20: Rubric for Factor 1.4 — Poisoning Detection

Level	Rating	Criteria
1	Non-Existent	No controls for detecting data poisoning. No awareness of poisoning as an attack vector.
2	Ad Hoc	Awareness of poisoning risks exists but detection is reactive and informal. No systematic monitoring of training data integrity.

Level	Rating	Criteria
3	Defined	Documented procedures for data integrity verification. Basic anomaly detection on training inputs. Pipeline integrity checks implemented but may not cover all data sources.
4	Managed	Automated anomaly detection across training data pipelines. Statistical monitoring for distribution anomalies. Documented incident response for suspected poisoning. Regular pipeline integrity audits.
5	Optimized	Advanced poisoning detection with ML-based anomaly identification. Real-time monitoring of all data pipelines. Automated quarantine of suspected poisoned data. Provenance verification at every pipeline stage. Threat intelligence integration for emerging poisoning techniques.

Table 21: Rubric for Factor 1.5 — Drift Detection and Monitoring

Level	Rating	Criteria
1	Non-Existent	No monitoring for model drift. No awareness of drift as a risk to production AI systems.
2	Ad Hoc	Some informal monitoring exists (e.g., periodic manual checks) but no systematic drift detection infrastructure. Alerting is reactive.
3	Defined	Drift monitoring infrastructure deployed with defined metrics (e.g., PSI, KL divergence). Alerting thresholds configured. Documented procedures for drift remediation.
4	Managed	Comprehensive drift monitoring across all production models. Automated alerting with defined escalation paths. Regular review cadence for monitoring thresholds. Retraining triggers documented and tested.
5	Optimized	Continuous drift monitoring with automated retraining pipelines. Predictive drift detection using leading indicators. Dynamic threshold adjustment. Automated model performance reporting. Integrated with model governance and stakeholder notification.

A.3 Domain 2: Output Liability and Content Risk

Table 22: Rubric for Factor 2.1 — Human-in-the-Loop Oversight

Level	Rating	Criteria
1	Non-Existent	No human oversight of AI outputs. AI systems operate autonomously without escalation mechanisms.
2	Ad Hoc	Some human review occurs but is informal, inconsistent, or limited to complaints. No defined criteria for what constitutes a 'high-stakes' decision requiring human review.
3	Defined	Documented policy defining high-stakes AI decisions requiring human review. Escalation triggers defined. Override mechanisms exist. Implementation may be inconsistent across all AI systems.

Level	Rating	Criteria
4	Managed	Consistent human oversight program with risk-calibrated review tiers. Automated escalation routing based on defined triggers. Override decisions documented and auditable. Regular review of escalation threshold appropriateness.
5	Optimized	Comprehensive oversight architecture with dynamic risk calibration. Automated confidence scoring with real-time escalation. Full audit trail of human review decisions. Regular effectiveness assessment of oversight program. Feedback loops from oversight findings to model improvement.

Table 23: Rubric for Factor 2.2 — Output Validation

Level	Rating	Criteria
1	Non-Existent	No validation of AI outputs before delivery to end users or business processes.
2	Ad Hoc	Some informal output review occurs but is inconsistent and undocumented. No systematic hallucination detection or fact-checking.
3	Defined	Documented output validation procedures. Basic fact-checking or accuracy verification processes in place. Hallucination detection awareness exists. Quality thresholds defined.
4	Managed	Automated validation pipelines with defined accuracy thresholds. Hallucination detection systems deployed. Confidence scoring applied to outputs. Regular accuracy measurement and reporting.
5	Optimized	Comprehensive automated validation with multiple verification layers. Real-time hallucination detection with continuous improvement. Accuracy metrics tracked against defined targets. Validation findings feed back to model training. Industry-leading output quality assurance.

Table 24: Rubric for Factor 2.3 — Disclaimer and Disclosure

Level	Rating	Criteria
1	Non-Existent	No disclosure of AI involvement in outputs or decisions. No transparency framework.
2	Ad Hoc	Some disclosure exists but is inconsistent or generic. Not calibrated to context or audience. May not cover all AI-enabled products or services.
3	Defined	Documented disclosure framework. AI involvement disclosed in relevant contexts. Output limitations communicated. Disclosure language reviewed by legal counsel.
4	Managed	Context-specific disclosure calibrated to audience and risk level. Disclosure coverage verified across all AI-enabled products. Regular review of disclosure adequacy. Compliance with jurisdictional disclosure requirements.
5	Optimized	Comprehensive transparency program with dynamic disclosure. Real-time disclosure calibration based on output confidence and risk. Proactive disclosure beyond regulatory minimums. Regular external assessment of transparency practices.

Table 25: Rubric for Factor 2.4 — Error Correction and Recall

Level	Rating	Criteria
1	Non-Existent	No procedures for correcting erroneous AI outputs or notifying affected parties.
2	Ad Hoc	Some informal correction processes exist but are reactive and undocumented. No systematic recall capability. Affected party notification is ad hoc.
3	Defined	Documented error correction procedures. Recall capability exists for AI outputs. Notification protocols defined for affected parties. Root cause analysis conducted for significant errors.
4	Managed	Automated error detection with defined response time objectives. Tested recall procedures. Structured notification workflows. Root cause analysis with documented corrective actions. Metrics tracked.
5	Optimized	Comprehensive error management with real-time detection and automated correction pipelines. Proactive error identification. Rapid recall capability tested regularly. Full traceability of outputs to affected parties. Continuous improvement from error analysis.

Table 26: Rubric for Factor 2.5 — Liability Documentation

Level	Rating	Criteria
1	Non-Existent	No contractual framework addressing AI output liability. Terms of service do not reference AI.
2	Ad Hoc	Some AI-related terms exist but are generic, incomplete, or not reviewed by legal counsel. Liability allocation unclear.
3	Defined	Documented liability allocation in terms of service and customer agreements. AI-specific provisions in vendor contracts. Legal review conducted. Coverage of major AI-enabled products.
4	Managed	Comprehensive contractual framework with jurisdiction-specific compliance. Regular legal review cadence. Indemnification provisions. Liability allocation documented for all AI-enabled products and services.
5	Optimized	Industry-leading contractual framework with dynamic updates for regulatory changes. Proactive liability risk assessment for new AI deployments. Cross-jurisdictional compliance. Insurance-aligned liability language. Regular external legal audit.

A.4 Domain 3: AI Supply Chain Security

Table 27: Rubric for Factor 3.1 — Vendor Due Diligence

Level	Rating	Criteria
1	Non-Existent	No formal vendor assessment for AI providers. Vendors selected without security evaluation.
2	Ad Hoc	Some informal vendor review occurs but is inconsistent. No standardized evaluation criteria. No ongoing monitoring.
3	Defined	Documented vendor assessment process with defined evaluation criteria. Security questionnaires used. Assessment conducted prior to engagement. Results documented.
4	Managed	Comprehensive vendor evaluation with risk-rated vendor tiers. Contractual security requirements. Ongoing vendor monitoring. Regular reassessment cadence. Findings tracked to remediation.
5	Optimized	Continuous vendor risk monitoring with automated intelligence feeds. Dynamic risk scoring. Vendor security integrated into procurement workflows. Regular on-site assessments for critical vendors. Vendor risk reporting to leadership.

Table 28: Rubric for Factor 3.2 — Dependency Mapping

Level	Rating	Criteria
1	Non-Existent	No inventory of AI dependencies. Third-party models, APIs, and services not documented.
2	Ad Hoc	Partial inventory exists but is incomplete or outdated. Dependencies known informally but not systematically documented.
3	Defined	Documented inventory of third-party AI components with dependency relationships. Inventory maintained and periodically updated. Impact assessment for major dependencies.
4	Managed	Comprehensive dependency inventory with automated discovery. Real-time updates. Impact analysis capabilities for supply chain changes. Concentration metrics calculated.
5	Optimized	Fully automated dependency mapping with real-time inventory. Predictive impact modeling for supply chain disruptions. Automated alerting on dependency changes. Integration with business continuity planning. Industry-leading supply chain visibility.

Table 29: Rubric for Factor 3.3 — Sub-processor Auditing

Level	Rating	Criteria
1	Non-Existent	No awareness of sub-processors in the AI pipeline. No audit rights or processes.
2	Ad Hoc	Some sub-processors identified but not systematically inventoried. No audit rights negotiated. No audit activity.

Level	Rating	Criteria
3	Defined	Sub-processor inventory maintained. Contractual audit rights established. Periodic audits conducted. Findings documented.
4	Managed	Comprehensive sub-processor register with regular audit execution. Audit findings tracked to remediation. Continuous compliance monitoring. Risk-rated audit frequency.
5	Optimized	Automated sub-processor monitoring with continuous compliance verification. Real-time audit capability. Proactive identification of sub-processor changes. Industry-leading supply chain audit program.

Table 30: Rubric for Factor 3.4 — Contractual Liability (Supply Chain)

Level	Rating	Criteria
1	Non-Existent	No AI-specific contractual provisions with vendors. Standard procurement terms only.
2	Ad Hoc	Some AI-related terms negotiated but inconsistently. Limited indemnification. No standardized AI vendor contract template.
3	Defined	AI-specific contractual provisions including indemnification, SLAs, and breach notification. Legal review of AI vendor contracts. Coverage of major vendors.
4	Managed	Standardized AI vendor contract template with comprehensive security provisions. Regular legal review. All AI vendors covered. SLAs with defined remedies. Breach notification requirements.
5	Optimized	Industry-leading contractual framework with dynamic updates. Cross-vendor security standards. Automated compliance monitoring against contractual obligations. Insurance-aligned liability provisions. Regular external review.

Table 31: Rubric for Factor 3.5 — Data Provenance Verification (Supply Chain)

Level	Rating	Criteria
1	Non-Existent	No verification of data provenance across the AI supply chain.
2	Ad Hoc	Some awareness of upstream data provenance but no systematic verification. Reliance on vendor representations without validation.
3	Defined	Documented verification procedures for upstream data provenance. Licensing compliance tracked for third-party data. Contractual provenance requirements in vendor agreements.
4	Managed	Automated provenance verification across the supply chain. Contractual provenance warranties from vendors. Regular audits of upstream data compliance. Remediation process for identified gaps.

Level	Rating	Criteria
5	Optimized	Comprehensive supply chain provenance assurance with real-time verification. End-to-end provenance tracking from raw data to model deployment. Automated compliance monitoring. Industry-leading data supply chain governance.

A.5 Domain 4: Regulatory and Compliance Alignment

Table 32: Rubric for Factor 4.1 — EU AI Act Readiness

Level	Rating	Criteria
1	Non-Existent	No awareness of or preparation for EU AI Act requirements.
2	Ad Hoc	General awareness of EU AI Act exists but no systematic assessment. No classification of AI systems against risk categories.
3	Defined	AI systems classified against EU AI Act risk categories. Compliance gap analysis conducted. Documented roadmap for addressing high-risk requirements.
4	Managed	Proactive compliance program with conformity assessment readiness. Technical documentation aligned with EU AI Act requirements. Human oversight mechanisms implemented. Regular compliance reviews.
5	Optimized	Full EU AI Act compliance with continuous monitoring for regulatory updates. Quality management system aligned with high-risk requirements. Proactive engagement with regulatory guidance. Compliance demonstrated through independent assessment.

Table 33: Rubric for Factor 4.2 — U.S. State Law Compliance

Level	Rating	Criteria
1	Non-Existent	No awareness of U.S. state-level AI regulations.
2	Ad Hoc	General awareness of some state regulations but no systematic mapping. Compliance posture unknown for most jurisdictions.
3	Defined	Compliance mapping conducted across applicable jurisdictions. Gap analysis documented. Remediation plans established for identified gaps.
4	Managed	Comprehensive jurisdiction-specific compliance programs. Continuous regulatory monitoring. Legal counsel review of AI deployments for state law compliance. Regular compliance assessments.
5	Optimized	Industry-leading multi-jurisdictional compliance with automated regulatory monitoring. Proactive compliance preparation for emerging state regulations. Legal review integrated into AI deployment process. External compliance audits.

Table 34: Rubric for Factor 4.3 — Cross-Border Data Transfer

Level	Rating	Criteria
1	Non-Existent	No protocols for cross-border data transfer compliance in AI contexts.
2	Ad Hoc	Awareness of cross-border requirements but no systematic compliance. Data transfer mechanisms not documented.
3	Defined	Documented transfer mechanisms (SCCs, adequacy decisions). GDPR compliance for EU data transfers. Data localization requirements identified.
4	Managed	Comprehensive transfer compliance with jurisdiction-specific mechanisms. Transfer impact assessments conducted. Automated compliance monitoring. Regular legal review.
5	Optimized	Industry-leading cross-border compliance with real-time transfer monitoring. Automated jurisdiction-specific data routing. Dynamic adaptation to regulatory changes. Comprehensive TIA program.

Table 35: Rubric for Factor 4.4 — Bias Testing and Fairness

Level	Rating	Criteria
1	Non-Existent	No algorithmic bias testing. No fairness metrics defined.
2	Ad Hoc	Some informal awareness of bias risk but no systematic testing. No defined fairness metrics.
3	Defined	Documented bias testing program with defined methodology. Fairness metrics selected. Testing conducted periodically. Results documented.
4	Managed	Regular bias testing with multiple fairness metrics. Remediation processes for identified biases. Testing integrated into model development lifecycle. Disparate impact analysis conducted.
5	Optimized	Continuous bias monitoring with automated fairness assessment. Independent fairness audits. Proactive bias prevention in model development. Industry-leading algorithmic fairness program. Public transparency reporting.

Table 36: Rubric for Factor 4.5 — Regulatory Reporting

Level	Rating	Criteria
1	Non-Existent	No procedures for AI-related regulatory reporting. No awareness of reporting obligations.
2	Ad Hoc	Some awareness of reporting obligations but no documented procedures. Reporting is reactive and ad hoc.
3	Defined	Documented regulatory reporting procedures. Notification timelines defined. Records retention policies in place. Legal counsel coordination established.

Level	Rating	Criteria
4	Managed	Comprehensive reporting program with jurisdiction-specific templates. Automated reporting workflows where applicable. Regular testing of notification procedures. Compliance with all applicable notification timelines.
5	Optimized	Industry-leading regulatory reporting with automated compliance monitoring. Pre-approved notification language. Real-time regulatory intelligence for emerging reporting requirements. Comprehensive records management. Proactive regulatory engagement.

A.6 Domain 5: Systemic Resilience and Continuity

Table 37: Rubric for Factor 5.1 — Model Diversification

Level	Rating	Criteria
1	Non-Existent	Complete dependence on a single AI model or provider. No diversification strategy.
2	Ad Hoc	Awareness of concentration risk but no active diversification. Informal consideration of alternatives.
3	Defined	Documented diversification strategy. Alternative models or providers identified. Fallback options documented but may not be tested.
4	Managed	Active multi-vendor or multi-model deployment. Tested failover procedures. Documented diversification rationale. Regular diversification reviews.
5	Optimized	Comprehensive diversification with automated failover. Real-time performance monitoring across providers. Dynamic load balancing. Regular diversification reviews with concentration risk metrics. Industry-leading model resilience.

Table 38: Rubric for Factor 5.2 — Failure Isolation

Level	Rating	Criteria
1	Non-Existent	No architectural controls for containing AI failures. Failures can propagate unchecked.
2	Ad Hoc	Some informal isolation exists but is not designed or documented. No circuit breakers or graceful degradation patterns.
3	Defined	Documented failure isolation architecture. Circuit breakers implemented for critical AI systems. Graceful degradation modes defined. Blast radius analysis conducted.
4	Managed	Comprehensive failure isolation with automated circuit breakers. Tested degradation modes. Documented blast radius boundaries. Regular failure injection testing.
5	Optimized	Industry-leading failure isolation with chaos engineering practices. Automated blast radius limitation. Dynamic circuit breaker calibration. Continuous failure resilience validation. Zero-trust architecture for AI components.

Table 39: Rubric for Factor 5.3 — Business Continuity

Level	Rating	Criteria
1	Non-Existent	No business continuity planning for AI system failures. No manual fallback procedures.
2	Ad Hoc	General BCP exists but does not address AI-specific scenarios. Manual fallbacks informal or untested.
3	Defined	AI-specific BCP documentation exists. Manual fallback procedures documented. Recovery time objectives defined. BCP not regularly tested.
4	Managed	Comprehensive AI-specific BCP with tested recovery procedures. Regular BCP exercises including AI failure scenarios. Validated manual fallback capabilities. Documented recovery playbooks.
5	Optimized	Industry-leading BCP with automated recovery orchestration. Regular AI-specific tabletop and functional exercises. Continuous BCP improvement from exercise findings. Recovery capability validated under realistic stress conditions.

Table 40: Rubric for Factor 5.4 — Concentration Monitoring

Level	Rating	Criteria
1	Non-Existent	No monitoring of AI infrastructure concentration. Single points of failure unknown.
2	Ad Hoc	Some informal awareness of dependencies but no systematic monitoring. Concentration risk not quantified.
3	Defined	Concentration risk identified and documented. Monitoring of key dependencies. Single points of failure cataloged. Basic concentration metrics defined.
4	Managed	Real-time concentration monitoring dashboards. Automated alerting on threshold breaches. Regular concentration risk reviews. Mitigation strategies documented for identified concentrations.
5	Optimized	Comprehensive concentration monitoring with predictive analytics. Dynamic threshold adjustment. Automated mitigation recommendations. Industry-leading concentration risk management. Integration with enterprise risk framework.

Table 41: Rubric for Factor 5.5 — AI Incident Response

Level	Rating	Criteria
1	Non-Existent	No incident response capability for AI-specific incidents. No playbooks or procedures.
2	Ad Hoc	General IT incident response may capture AI incidents informally. No AI-specific playbooks. No model quarantine or output recall capability.
3	Defined	AI-specific incident response playbooks documented. Model quarantine procedures defined. Output recall capability exists. Stakeholder notification protocols established.

Level	Rating	Criteria
4	Managed	Tested AI IR playbooks with regular tabletop exercises. Automated quarantine capabilities. Documented response time objectives. Post-incident review processes. Lessons learned integration.
5	Optimized	Industry-leading AI IR with automated detection, quarantine, and notification. Regular functional exercises. Continuous improvement from incident analysis. Proactive threat hunting for AI-specific attack patterns. Cross-organizational coordination capability.

Annex B: Regulatory Framework Crosswalks (Informative)

This annex maps AIRS factors to three major regulatory frameworks: NIST AI RMF 1.0, ISO/IEC 42001:2023, and the EU Artificial Intelligence Act. These mappings are informative and intended to assist organizations in leveraging existing compliance investments for AIRS assessment. The mappings identify the primary alignment points; organizations should consult the source frameworks for complete requirements.

B.1 AIRS and NIST AI RMF 1.0

Table B.1: AIRS to NIST AI RMF 1.0 Crosswalk

AIRS Factor	NIST AI RMF Functions	Primary Subcategories
1.1 Training Data Provenance	GOVERN, MAP	GOVERN 6.1, MAP 2.3
1.2 Adversarial Robustness	GOVERN, MEASURE	GOVERN 4.3, MEASURE 2.6, 2.7
1.3 Model Versioning	GOVERN, MANAGE	GOVERN 1.5, MANAGE 2.2
1.4 Poisoning Detection	MEASURE, MANAGE	MEASURE 2.7, MANAGE 2.4
1.5 Drift Detection	MEASURE, MANAGE	MEASURE 2.5, 2.6, MANAGE 2.2
2.1 Human-in-the-Loop	GOVERN, MANAGE	GOVERN 1.4, MANAGE 3.2
2.2 Output Validation	MEASURE	MEASURE 2.3, 2.5, 2.11
2.3 Disclaimer/Disclosure	MAP, GOVERN	MAP 5.2, GOVERN 4.2
2.4 Error Correction	MANAGE	MANAGE 2.4, 4.1
2.5 Liability Documentation	GOVERN	GOVERN 6.1, 6.2
3.1 Vendor Due Diligence	GOVERN, MAP	GOVERN 6.1, MAP 5.1
3.2 Dependency Mapping	MAP	MAP 1.5, 1.6, 3.4
3.3 Sub-processor Auditing	GOVERN	GOVERN 6.1, 6.2
3.4 Contractual Liability	GOVERN	GOVERN 6.1, 6.2
3.5 Supply Chain Provenance	GOVERN, MAP	GOVERN 6.1, MAP 2.3
4.1 EU AI Act Readiness	GOVERN, MAP	GOVERN 1.1, MAP 1.1, 1.2
4.2 U.S. State Compliance	GOVERN	GOVERN 1.1, 1.2
4.3 Cross-Border Transfer	GOVERN	GOVERN 6.1

AIRS Factor	NIST AI RMF Functions	Primary Subcategories
4.4 Bias Testing	MEASURE	MEASURE 2.6, 2.10, 2.11
4.5 Regulatory Reporting	GOVERN, MANAGE	GOVERN 5.1, MANAGE 4.1
5.1 Model Diversification	MANAGE, MAP	MANAGE 2.1, MAP 3.5
5.2 Failure Isolation	MANAGE	MANAGE 2.1, 2.3
5.3 Business Continuity	MANAGE	MANAGE 2.1, 4.2
5.4 Concentration Monitoring	MAP, MEASURE	MAP 3.4, MEASURE 3.2
5.5 AI Incident Response	MANAGE	MANAGE 4.1, 4.2

B.2 AIRS and ISO/IEC 42001:2023

Table B.2: AIRS to ISO/IEC 42001:2023 Crosswalk

AIRS Domain	ISO 42001 Clause(s)	Alignment Notes
D1: Model Integrity	8.2, 8.4, A.6, A.7	Model lifecycle controls, data quality, robustness testing
D2: Output Liability	8.3, 8.4, A.8, A.9	Output management, human oversight, transparency
D3: Supply Chain	A.10, 7.1.3, 8.1	Third-party risk, procurement, operational controls
D4: Regulatory	4.1, 4.2, 6.1, 9.1	Context, interested parties, risk, monitoring
D5: Resilience	8.1, A.5, A.6	Operational planning, availability, continuity

B.3 AIRS and EU Artificial Intelligence Act

Table B.3: AIRS to EU AI Act Crosswalk

AIRS Domain	EU AI Act Article(s)	Requirements Addressed
D1: Model Integrity	Art. 9, 10, 15	Risk management, data governance, accuracy/robustness
D2: Output Liability	Art. 13, 14, 52	Transparency, human oversight, disclosure obligations
D3: Supply Chain	Art. 25, 28	Distributor obligations, deployer obligations
D4: Regulatory	Art. 6, 43, 49, 62	Classification, conformity assessment, registration, reporting
D5: Resilience	Art. 9, 15	Risk management systems, resilience measures

Annex C: Implementation Guidance (Informative)

C.1 First-Time Implementation

Organizations adopting AIRS for the first time should consider the following implementation approach:

- Inventory AI systems. Catalog all AI systems in production, development, and procurement, including third-party AI components. This inventory forms the assessment boundary.
- Conduct a gap assessment. Using the maturity rubrics in Annex A, perform an informal self-assessment to identify current maturity levels and priority gaps.
- Prioritize remediation. Focus on factors with the greatest impact on composite score improvement. Domain 1 factors carry the highest weight and often yield the largest score improvements.
- Build evidence. Establish documentation practices that produce the evidence categories described in Section 8.3. Many organizations underinvest in documentation of existing practices.
- Engage assessors. For the initial formal assessment, consider engaging an independent assessor to establish a credible baseline and identify blind spots.

C.2 Integration with Existing Risk Frameworks

Organizations with existing ISO 27001, ISO 42001, or NIST CSF programs can leverage significant existing evidence for AIRS assessment:

- ISO 27001 certified organizations will find that Annex A controls addressing access management, asset management, and supplier relationships provide partial evidence for AIRS Domains 3 and 5.
- ISO 42001 certified organizations will find substantial overlap with AIRS across all five domains, particularly in model lifecycle management, data governance, and human oversight.
- NIST CSF adopters will find that Identify and Protect functions align with AIRS Domains 1 and 3, while Detect and Respond align with Domains 2 and 5.

C.3 Guidance for Carriers

Insurance carriers integrating AIRS into underwriting workflows should consider:

- Establishing AIRS tier thresholds as minimum requirements for AI-related coverage endorsements
- Incorporating AIRS scores into renewal assessment alongside traditional cybersecurity ratings
- Using domain-level scores (not just composite) to identify specific coverage exclusions or sub-limits

- Offering premium incentives for entities that improve from Tier 3 to Tier 2, or Tier 2 to Tier 1
- Including AIRS reassessment requirements as policy conditions for conditional coverage

C.4 Guidance for Regulators

Regulators considering AIRS for incorporation into guidance or examination procedures should note that the standard:

- Is published as an open standard with no licensing restrictions on regulatory use
- Provides citation formats specifically designed for regulatory filings (Section 10.2)
- Distinguishes between internal insurer AI governance (out of scope) and underwritten entity AI risk (in scope)
- Aligns with existing NIST and ISO frameworks already referenced in regulatory guidance
- Defines conformance criteria that enable consistent regulatory examination

Annex D: Version History and Change Log (Informative)

D.1 Version History

Version	Date	Change Type	Summary
1.0	April 2026	Initial Publication	First release of AIRS v1.0 Open Standard Specification. Defines five scoring domains, twenty-five risk factors, composite scoring formula, four rating tiers, assessment process requirements, and conformance criteria.

D.2 Errata

No errata have been issued for this version.

Errata may be submitted to standards@aisecurityintelligence.com. Confirmed errata will be published in a subsequent document version update.

Annex E: NAIC Regulatory Adoption Guide (Informative)

E.1 Purpose and Scope

This annex provides guidance for state insurance departments, the National Association of Insurance Commissioners (NAIC), and market participants on the relationship between AIRS and the NAIC's existing AI regulatory framework. It maps AIRS scoring domains to NAIC instruments, clarifies complementary scope, and provides template regulatory language for adoption.

> This annex is informative. It does not modify the normative requirements of the AIRS standard and does not constitute legal or regulatory advice. State insurance departments should consult applicable state law and regulation when considering AIRS adoption.

E.2 The Complementary Scope Distinction

The NAIC's AI regulatory framework and AIRS address two fundamentally distinct risk vectors that together constitute the complete AI risk landscape facing the insurance industry. Understanding this distinction is essential for effective regulatory adoption.

Table E.1: Scope Comparison: NAIC AI Framework vs. AIRS

Dimension	NAIC AI Framework	AIRS
Primary question	How does the insurer use AI in its own operations?	How secure is the AI posture of entities the insurer underwrites?
Direction of evaluation	Inward — evaluates the carrier's internal AI governance	Outward — evaluates the AI security of insured entities
Regulatory instrument	Model Bulletin (Dec 2023); AI Systems Evaluation Tool (2026 pilot)	Open standard specification for underwriting risk assessment
What it measures	Governance, fairness, accountability, transparency of insurer AI	Security posture, vulnerability exposure, resilience of assessed entities
Risk type addressed	Operational and conduct risk from insurer's own AI use	Underwriting risk from insured entity's AI deployment
Adoption status	24 states adopted Model Bulletin; 12-state Evaluation Tool pilot (Mar–Sep 2026)	Published April 2026 as open standard

The NAIC framework asks: "Are our carriers governing their AI responsibly?" AIRS asks: "Are the entities our carriers insure managing their AI security risk?" These are complementary questions. A carrier may have exemplary internal AI governance while simultaneously underwriting entities with catastrophic AI security gaps. The NAIC framework addresses the first risk. AIRS addresses the second.

E.3 NAIC Instrument Mapping

This section maps the three primary NAIC AI regulatory instruments to AIRS scoring domains, identifying where they converge, where they complement, and where AIRS extends beyond existing NAIC coverage.

E.3.1 NAIC Principles of Artificial Intelligence (2020)

The NAIC adopted five AI principles in August 2020: fair and ethical use, accountability, compliance with laws and regulations, transparency, and safe, secure, fair, and robust systems. These principles guide insurer behavior. AIRS operationalizes equivalent principles as measurable factors applied to assessed entities.

Table E.2: NAIC AI Principles to AIRS Domain Mapping

NAIC Principle	NAIC Application	AIRS Counterpart	AIRS Domain(s)
Fair and Ethical	Insurer AI avoids unfair discrimination	Bias testing, human oversight of AI outputs	D2 (Factors 2.4, 2.1)
Accountability	Insurer governance structures for AI	Governance maturity, audit trails, version control	D1 (Factor 1.3), D4
Compliance	Insurer compliance with insurance law	Regulatory alignment across jurisdictions	D4 (Factors 4.1–4.5)
Transparency	Explainability of insurer AI decisions	Disclosure controls, content provenance	D2 (Factors 2.3, 2.5)
Safe, Secure, Robust	Reliability of insurer AI systems	Model integrity, supply chain security, resilience	D1, D3, D5

E.3.2 NAIC Model Bulletin on Use of AI by Insurers (December 2023)

The Model Bulletin, adopted by 24 states as of March 2025, requires insurers to develop written AI programs (AIS Programs) addressing governance, risk management, internal controls, and third-party oversight for AI used in regulated insurance practices. The Bulletin's four key sections map to AIRS as follows:

Table E.3: Model Bulletin Sections to AIRS Mapping

Bulletin Section	Requirement	AIRS Parallel	Relationship
1. AIS Program	Written AI program with governance, risk mgmt, internal audit	Domains 1–5 assess the substance that AIS Programs should produce	Complementary: Bulletin mandates the program; AIRS measures its effectiveness in assessed entities
2. Governance	Accountability structure, transparency, fairness in AI design	Domain 4 (Regulatory Alignment) + Domain 2 (Output Liability)	Converging: Both require governance; AIRS adds quantitative scoring

Bulletin Section	Requirement	AIRS Parallel	Relationship
3. Risk Management	Controls for data practices, validation, testing, privacy	Domain 1 (Model Integrity) + Domain 3 (Supply Chain)	Converging: Bulletin defines expectations; AIRS rubrics define maturity levels
4. Third-Party Oversight	Vendor management, due diligence on third-party AI	Domain 3 (AI Supply Chain Security), Factors 3.1–3.5	Extending: AIRS provides structured scoring for supply chain risk that the Bulletin references qualitatively

E.3.3 NAIC AI Systems Evaluation Tool (2026)

The AI Systems Evaluation Tool, currently in a 12-state pilot program (March–September 2026) with formal adoption expected at the NAIC Fall National Meeting in November 2026, provides four exhibits for regulators to assess insurer AI usage during market conduct and financial examinations. The Tool evaluates how insurers deploy AI internally. AIRS evaluates the AI security posture of entities those insurers underwrite.

Table E.4: Evaluation Tool Exhibits to AIRS Domain Mapping

Evaluation Tool Exhibit	Purpose (Insurer-Facing)	AIRS Analog (Entity-Facing)	Gap AIRS Fills
Exhibit A: Quantify AI Usage	Inventory insurer's AI systems across operations	AIRS assessment scoping identifies all AI systems in the assessed entity	AIRS applies this to underwritten entities, not just the carrier
Exhibit B: Governance Framework	Assess insurer's AI governance and risk management	Domains 1, 4: Model integrity governance + regulatory compliance	AIRS scores governance maturity on a 1–5 rubric scale, producing quantifiable output
Exhibit C: High-Risk AI Details	Detailed review of insurer's high-risk AI models	Domain 1: Model Integrity and Validation (all 5 factors)	AIRS evaluates model risk in entities being underwritten, not just carrier models
Exhibit D: AI Data Details	Data sources, types, and risk of adverse consumer impact	Domain 1 (Factor 1.1: Data Provenance) + Domain 2 (output controls)	AIRS extends data risk assessment to the underwriting target

E.4 Adoption Pathways for State Insurance Departments

State insurance departments may reference or incorporate AIRS through several regulatory mechanisms, depending on existing statutory authority and departmental practice.

Bulletin or guidance issuance. A state DOI issues a bulletin recommending or requiring that carriers consider AIRS-conformant assessments as part of AI-related underwriting. This is the lowest-friction adoption path and mirrors the mechanism used for Model Bulletin adoption.

Market conduct examination integration. State examiners incorporate AIRS scoring into market conduct reviews of carriers' AI underwriting practices. When reviewing how a carrier evaluates AI risk in its book, examiners reference AIRS as a benchmark for methodological adequacy.

Rate filing requirements. States may require or recommend that rate filings for AI-related coverage classes include AIRS-conformant risk assessments as supporting documentation, establishing AIRS as part of the actuarial basis for AI risk pricing.

Financial examination supplement. AIRS scoring data may supplement financial examination procedures when evaluating a carrier's reserve adequacy for AI-related loss exposure. AIRS tier classifications provide a structured basis for assessing concentration of underwriting risk.

E.5 Sample Regulatory Language

The following template language is provided for state insurance departments considering AIRS adoption. These templates are illustrative and should be adapted to conform with applicable state statutory authority and departmental drafting conventions.

E.5.1 Bulletin Language: Recommended Use

> "The Department recommends that insurers writing coverage for entities deploying artificial intelligence systems incorporate standardized AI security risk assessments into their underwriting processes. The AI Insurance Readiness Score (AIRS) v1.0 Open Standard Specification, published by AI Security Intelligence LLC, provides a transparent, reproducible methodology for evaluating AI security posture across five scoring domains. Insurers are encouraged to consider AIRS-conformant assessments, or assessments of equivalent rigor and scope, when evaluating AI-related risk exposure in their underwriting portfolios."

E.5.2 Rate Filing Language: Supporting Documentation

> "Rate filings for coverage classes with material AI-related loss exposure should include documentation of the methodology used to evaluate insured entities' AI security risk. Filings referencing the AI Insurance Readiness Score (AIRS) standard or AIRS-conformant assessments shall cite the applicable standard version and assessment date. The Department considers AIRS-conformant methodologies as an acceptable basis for AI risk differentiation in rate development."

E.5.3 Market Conduct Examination Language

> "When examining an insurer's underwriting practices for AI-related coverage, examiners may request documentation of the methodology used to evaluate AI security risk in insured entities. The AI Insurance Readiness Score (AIRS) v1.0 Open Standard Specification provides a reference framework for evaluating the adequacy and consistency of such methodologies. Examiners should assess whether the insurer's AI risk evaluation addresses, at minimum, the five AIRS scoring domains: model integrity, output liability, supply chain security, regulatory alignment, and systemic resilience."

E.6 Relationship to NAIC Evaluation Tool Pilot

The NAIC AI Systems Evaluation Tool pilot, operating across twelve participating states (California, Colorado, Connecticut, Florida, Iowa, Louisiana, Maryland, Pennsylvania, Rhode Island, Vermont, Virginia, and Wisconsin) from March through September 2026, is developing the methodology regulators will use to evaluate insurer AI governance. AIRS complements this effort by providing the methodology carriers can use to evaluate AI risk in the entities they insure.

As the Evaluation Tool matures toward formal adoption at the November 2026 NAIC Fall National Meeting, AIRS offers a ready framework for the underwriting-facing counterpart that the Tool does not address. Together, the NAIC Evaluation Tool and AIRS provide regulators with visibility into both sides of the AI risk equation: how carriers govern their own AI, and how carriers assess AI risk in their underwriting portfolios.

E.7 State Adoption Tracking

As of the publication date of this specification, the following table summarizes NAIC AI framework adoption by state. AIRS adoption status will be maintained at aisecurityintelligence.com/airs-adoption.

Table E.5: NAIC AI Framework State Adoption Summary (as of April 2026)

NAIC Instrument	Adoption Status	States
Model Bulletin (Dec 2023)	24 states adopted	AK, AR, CT, DE, DC, IL, IA, KY, MD, MA, MI, NE, NV, NH, NJ, NC, OK, PA, RI, VT, VA, WA, WV, WI
AI Evaluation Tool Pilot (2026)	12 states participating	CA, CO, CT, FL, IA, LA, MD, PA, RI, VT, VA, WI
AIRS v1.0 Standard	Published April 2026	Adoption tracking at aisecurityintelligence.com/airs-adoption

NOTE: The state adoption list for the NAIC Model Bulletin is current as of March 2025. Additional states may have adopted subsequent to this date. Consult the NAIC and individual state DOI websites for the most current adoption status.

References

[1] National Institute of Standards and Technology. (2023). Artificial Intelligence Risk Management Framework (AI RMF 1.0). NIST AI 100-1. https://doi.org/10.6028/NIST.AI.100-1

[2] International Organization for Standardization. (2023). Information technology — Artificial intelligence — Artificial intelligence management system. ISO/IEC 42001:2023.

[3] European Parliament and Council. (2024). Regulation (EU) 2024/1689 laying down harmonised rules on artificial intelligence (Artificial Intelligence Act). Official Journal of the European Union.

[4] National Institute of Standards and Technology. (2020). Security and Privacy Controls for Information Systems and Organizations. NIST SP 800-53, Revision 5. https://doi.org/10.6028/NIST.SP.800-53r5

[5] FIRST.Org, Inc. (2023). Common Vulnerability Scoring System version 4.0: Specification Document. https://www.first.org/cvss/v4.0/specification-document

[6] FAIR Institute. (2025). Factor Analysis of Information Risk (FAIR) Standard, Version 3.0. https://www.fairinstitute.org/

[7] International Organization for Standardization. (2022). Information technology — Artificial intelligence — Artificial intelligence concepts and terminology. ISO/IEC 22989:2022.

[8] National Association of Insurance Commissioners. (2023). Model Bulletin on the Use of Artificial Intelligence Systems by Insurers. https://content.naic.org/

[9] National Institute of Standards and Technology. (2024). Cybersecurity Framework 2.0. NIST CSWP 29. https://doi.org/10.6028/NIST.CSWP.29

[10] A.M. Best Rating Services. (2024). Best's Credit Rating Methodology. Version 011624.

[11] National Association of Insurance Commissioners. (2020). Principles of Artificial Intelligence. Adopted August 2020. https://content.naic.org/insurance-topics/artificial-intelligence

[12] National Association of Insurance Commissioners. (2026). AI Systems Evaluation Tool, Version 4.0. Big Data and Artificial Intelligence (H) Working Group. Pilot program March–September 2026.

Glossary

This glossary provides brief definitions for quick reference. Formal definitions for terms used normatively in this specification appear in Section 3.

Abbreviation	Definition
AIRS	AI Insurance Readiness Score — the composite risk score defined by this standard
API	Application Programming Interface
BCP	Business Continuity Plan
CCPA	California Consumer Privacy Act
CI/CD	Continuous Integration / Continuous Deployment
CVSS	Common Vulnerability Scoring System
EU AI Act	European Union Artificial Intelligence Act (Regulation (EU) 2024/1689)
FAIR	Factor Analysis of Information Risk
GDPR	General Data Protection Regulation (Regulation (EU) 2016/679)
GWP	Gross Written Premium
IR	Incident Response
ISO	International Organization for Standardization
KL Divergence	Kullback-Leibler Divergence — a statistical measure of distribution difference
LLM	Large Language Model
ML	Machine Learning
NAIC	National Association of Insurance Commissioners
NIST	National Institute of Standards and Technology
NIST AI RMF	NIST Artificial Intelligence Risk Management Framework
PSI	Population Stability Index — a measure of distribution shift
SCC	Standard Contractual Clauses
SLA	Service Level Agreement
SoA	Statement of Applicability
TIA	Transfer Impact Assessment