AI SECURITY INTELLIGENCE

ASI-VSS

Vulnerability Scoring Standard

Methodology & Empirical Foundation

Version 1.0 | April 2026

EXECUTIVE SUMMARY

ASI-VSS: Vulnerability Scoring Standard v1.0

Artificial intelligence vulnerabilities are outpacing the security industry's ability to score them. Traditional frameworks—CVSS v4.0, the de facto standard for two decades—were built for a different era: deterministic software with well-defined boundaries. They were not designed to assess AI-specific attack vectors such as autonomous agent compromise, multi-agent propagation, model poisoning persistence, or cascading failures across interconnected systems. The result is a growing gap between actual AI risk and the industry's capacity to quantify it.

OWASP took an important step in March 2026 with the pre-release of AIVSS v0.8, introducing 10 agentic amplification factors—a welcome advance for the field. As a pre-release framework, AIVSS v0.8 naturally has room to evolve: its ternary scoring model (0 / 0.5 / 1) limits granularity, propagation modeling and empirical incident grounding remain on the roadmap, and insurance underwriting integration is not addressed. Its v1.0 public review opens April 16, 2026.

ASI-VSS builds on this foundation. Developed by AI Security Intelligence LLC, the Vulnerability Scoring Standard v1.0 introduces a five-dimension, 0–100 scoring architecture grounded in 174 analyst-verified AI security incidents collected between January 2025 and March 2026. The standard classifies vulnerabilities across 17 empirically-derived categories, maps each to MITRE ATLAS, NIST AI RMF, and OWASP risk taxonomies, and provides direct insurance loss category scoring—extending the field into territory no existing standard yet covers.

> KEY FINDING
>
> The highest-scoring vulnerability categories in our dataset are the AI-native ones that existing frameworks were not built to differentiate: Agent Compromise (avg 80.5), Cascading Agent Failures (avg 80.3), and Memory & Context Poisoning (avg 75.0). These categories require dedicated scoring dimensions for AI amplification, propagation modeling, and multi-agent dynamics—the precise dimensions ASI-VSS introduces.

ASI-VSS v1.0 publishes April 15, 2026. This document presents the complete methodology, the empirical foundation underlying our scoring calibrations, and implementation guidance for practitioners, carriers, and regulators seeking to integrate rigorous AI vulnerability scoring into their security programs.

PART I

The Scoring Gap

Chapter 1

The Evolution of Vulnerability Scoring

The Common Vulnerability Scoring System (CVSS) has served as the backbone of vulnerability prioritization since 2005—a remarkable two-decade contribution to the security community. Version 4.0, released in November 2023, further refined its metrics for traditional software and network vulnerabilities. CVSS was architected for a world of deterministic software systems with well-defined boundaries—a world that the rise of autonomous AI agents, multi-model architectures, and interconnected AI systems has fundamentally expanded.

Where CVSS v4.0 Meets Its Design Boundary

CVSS v4.0 excels at what it was designed to do. Three dimensions of AI-specific risk, however, fall outside its original scope:

AI Amplification. When an autonomous agent is compromised, the blast radius is not bounded by the agent's own privileges—it extends to every tool, API, and downstream system the agent can reach. CVSS's current model does not account for autonomy amplification, tool reach, or non-deterministic exploitation paths.

Propagation modeling. AI systems exhibit propagation characteristics unlike traditional software. A poisoned training dataset persists through model retraining cycles. A compromised agent in a multi-agent network can corrupt downstream agents within seconds. CVSS's event-based model evaluates vulnerabilities individually rather than as propagation phenomena.

Insurance mapping. Cyber insurance underwriting requires mapping vulnerabilities to specific coverage triggers—business interruption, cyber liability, D&O, data breach notification. CVSS was designed to produce a technical severity score—mapping that score to financial or insurance impact categories requires a separate layer.

OWASP AIVSS v0.8 Assessment

The OWASP AI Vulnerability Scoring System v0.8, released March 19, 2026, represents a significant milestone—the first community-driven effort to extend vulnerability scoring to AI-specific factors. Our intelligence team reviewed AIVSS v0.8 and sees it as an important early contribution to a conversation the industry urgently needs.

AIVSS v0.8 introduces 10 agentic amplification factors—a meaningful conceptual advance. As with any pre-release standard, there are areas where subsequent versions may choose to expand:

• Scoring granularity: Ternary factors (0, 0.5, or 1) provide a starting framework; continuous scoring could offer finer differentiation between materially different risk levels
• Propagation modeling: Dedicated dimensions for persistence duration, propagation velocity, cross-system reach, and detection difficulty are not included
• Insurance integration: Mapping vulnerability severity to underwriting criteria and coverage triggers is not addressed
• Empirical grounding: Scoring weights draw from expert surveys and 1,900 public comments—a strong consensus base, but doesn't include key observed incident data
• Taxonomy coverage: 10 risk categories provide coverage, but emerging categories such as Regulatory Violations, Adversarial Evasion, and Shadow AI warrant inclusion as the threat landscape evolves

The Scoring Compression Problem

Both CVSS and OWASP AIVSS use a 0–10 scale. For traditional software vulnerabilities, this provides adequate differentiation. For AI vulnerabilities—which span a wider severity spectrum due to amplification and propagation dynamics—a 0–10 scale compresses meaningfully different risks into the same integer band. An agent compromise scoring 8.5 and a cascading multi-agent failure scoring 8.7 appear nearly identical on a 0–10 scale, yet represent fundamentally different risk profiles requiring different response protocols.

Comparative Analysis

Dimension	CVSS v4.0	OWASP AIVSS v0.8	ASI-VSS v1.0
Scale	0–10	0–10	0–100 (10× granularity)
AI-Specific Scoring	Not in scope	10 agentic factors (ternary: 0/0.5/1)	5 AI sub-factors (continuous scoring)
Propagation Modeling	Not in scope	Not addressed	Dedicated D3 dimension
Insurance Mapping	Not in scope	Not addressed	D4 insurance loss categories
Empirical Basis	Expert consensus	Expert surveys + 1,900 comments	174 analyst-verified real-world incidents
Taxonomy	N/A (individual vulns)	10 risk categories	17 empirically-derived categories
Framework Mapping	CVE/NVD	MAESTRO/NIST indirect	MITRE ATLAS + NIST AI RMF + OWASP direct

Dimension	CVSS v4.0	OWASP AIVSS v0.8	ASI-VSS v1.0
Methodology Rigor	Mature (v4.0, 20 years)	Pre-release (v0.8)	Empirically-grounded, reproducible

Table 1.1 — Comparative analysis of vulnerability scoring standards

PART II

The ASI-VSS Architecture

Chapter 2

Five-Dimension Scoring Model

ASI-VSS employs a five-dimension scoring architecture that produces a composite score on a 0–100 scale. Each dimension contributes equally (0–20 points), reflecting the principle that technical severity, AI amplification, propagation dynamics, enterprise impact, and exploitability are co-equal determinants of vulnerability severity.

```
ASI-VSS Score = D1 + D2 + D3 + D4 + D5    (range: 0-100)
```

Dimension	Name	Range	Focus
D1	Technical Exploit Severity	0–20	CIA triad + attack complexity
D2	AI Capability Amplification	0–20	Autonomy, non-determinism, tool reach, propagation
D3	Propagation & Persistence	0–20	Temporal & spatial spread characteristics
D4	Enterprise & Financial Impact	0–20	Financial loss, regulatory, insurance mapping
D5	Exposure & Exploitability	0–20	Threat landscape context & exploit maturity

Table 2.1 — ASI-VSS five-dimension architecture overview

D1: Technical Exploit Severity (0–20)

Measures the inherent technical severity of the vulnerability independent of AI context, aligning with the CIA triad familiar to practitioners.

Sub-Factor	Range	Description
Confidentiality Impact	0–5	Degree of unauthorized data access enabled
Integrity Impact	0–5	Degree of unauthorized data/model modification

Sub-Factor	Range	Description
Availability Impact	0–5	Degree of service disruption or denial
Attack Complexity (inv.)	0–5	Inverse of exploitation difficulty (easy = high)

Empirical average: D1 = 11.4 / 20 across 174 incidents

D2: AI Capability Amplification (0–20)

Measures how AI-specific capabilities amplify the base vulnerability. This is where ASI-VSS diverges fundamentally from CVSS—and where the most dangerous AI vulnerabilities concentrate.

Sub-Factor	Range	Description
Autonomy Amplification	0–4	Exploitation by/through autonomous agents
Non-Determinism Risk	0–4	Unpredictability from model non-determinism
Tool & API Reach	0–4	Blast radius expansion via agent tool access
Learning & Adaptation	0–4	Persistence through model updates/fine-tuning
Multi-Agent Propagation	0–4	Cascade potential across agent networks

Empirical average: D2 = 6.4 / 20 across 174 incidents

D3: Propagation & Persistence (0–20)

Measures the temporal and spatial spread characteristics unique to AI systems—a dimension entirely absent from both CVSS and OWASP AIVSS.

Sub-Factor	Range	Description
Persistence Duration	0–5	Ephemeral → permanent (training data, weights)
Propagation Velocity	0–5	Hours → minutes → seconds spread rate
Cross-System Reach	0–5	Number of systems/models/agents affected
Detection Difficulty	0–5	Difficulty of post-exploitation detection

Empirical average: D3 = 10.6 / 20 across 174 incidents

D4: Enterprise & Financial Impact (0–20)

Measures the organizational consequence of exploitation, including insurance-relevant loss categories. This dimension enables direct mapping from vulnerability severity to underwriting decisions.

Sub-Factor	Range	Description
Financial Loss Magnitude	0–5	Direct financial impact (operational loss, fraud, ransom)
Regulatory Exposure	0–5	Regulatory triggers (GDPR, EU AI Act, NIST, SOX)
Reputational Damage	0–5	Brand and trust impact severity
Insurance Loss Category	0–5	Maps to cyber insurance coverage triggers

Empirical average: D4 = 13.4 / 20 across 174 incidents

D5: Exposure & Exploitability (0–20)

Measures the current threat landscape context—market defense maturity, exploit availability, and the breadth of the vulnerable attack surface.

Sub-Factor	Range	Description
Exploit Maturity	0–5	Theoretical → PoC → weaponized → actively exploited
Vendor Defense Maturity	0–5	Market detection/prevention capability
Attack Surface Breadth	0–5	Deployment breadth of vulnerable technology
Intelligence Confidence	0–5	Quality of supporting evidence

Empirical average: D5 = 13.4 / 20 across 174 incidents

Why 0–100, Not 0–10

Granularity. AI vulnerabilities exhibit a wider severity spectrum than traditional software flaws. A 0–10 scale compresses meaningfully different vulnerabilities into the same integer band. The 0–100 scale allows precise differentiation—a score of 73 represents materially different risk than a score of 78.

Insurance compatibility. Actuarial calculations require finer granularity than a 0–10 scale provides. Insurance loss projections for a score-73 vulnerability differ meaningfully from those for a score-78 vulnerability—a distinction invisible on the CVSS scale.

Dimensional transparency. Five dimensions at 0–20 each allow stakeholders to see exactly where severity concentrates: a score of 75 composed as D1=18, D2=17, D3=16, D4=14, D5=10 tells a

fundamentally different story than D1=15, D2=15, D3=15, D4=15, D5=15.

CVSS Interoperability

For practitioners operating in CVSS-integrated environments, ASI-VSS provides a direct mapping:

```
CVSS_Equivalent = ASI-VSS_Score / 10    (range: 0.0-10.0)
```

ASI-VSS Score	CVSS Equivalent	Severity Band
0–19	0.0–1.9	LOW
20–39	2.0–3.9	MODERATE
40–59	4.0–5.9	ELEVATED
60–79	6.0–7.9	HIGH
80–100	8.0–10.0	CRITICAL

Table 2.2 — CVSS interoperability mapping

This mapping preserves the practitioner's existing CVSS-based workflows while providing the 10× granularity advantage of the ASI-VSS scale. A CVSS equivalent of 7.3 (ASI-VSS 73) and 7.8 (ASI-VSS 78) represent materially different risk profiles—a distinction invisible on the native CVSS scale.

Chapter 3

Severity Classifications & Insurance Mapping

ASI-VSS maps composite scores to a five-tier severity classification system. Each tier carries defined response timelines and insurance implications, enabling organizations to operationalize vulnerability scores directly into their incident response and risk management workflows.

Classification	Score Range	Response Timeline	Insurance Implication
CRITICAL	85–100	Immediate action required	Mandatory carrier notification
SEVERE	70–84	Action within 48 hours	Material risk event — assess coverage
ELEVATED	50–69	Action within 1 week	Monitored risk — log for renewal
MODERATE	30–49	Scheduled remediation	Standard risk management
LOW	0–29	Monitor and assess	No coverage implication

Table 3.1 — ASI-VSS severity classification system

D4 Insurance Loss Category Scoring Rubric

The Insurance Loss Category sub-factor within D4 provides direct mapping from vulnerability severity to standard cyber insurance coverage triggers:

Score	Coverage Trigger	Typical Claim Range
5	Business Interruption + Cyber Liability + D&O;	>$50M
4	Cyber Liability + Regulatory Defense	$10M–$50M
3	Cyber Liability + Data Breach Notification	$1M–$10M
2	First-Party Data Recovery	$100K–$1M
1	Incident Response Costs Only	<$100K
0	No insurable loss	—

Table 3.2 — Insurance loss category mapping

Chapter 4

Scoring Algorithm

ASI-VSS produces scores through a five-step algorithm that combines analyst assessment with empirical calibration, ensuring both accuracy and reproducibility.

Step 1: Raw Dimension Scoring

Each of the five dimensions (D1–D5) is scored independently using the sub-factor rubrics defined in Chapter 2. Sub-factor scores are summed to produce the dimension score (0–20). Analysts assess each sub-factor based on available evidence, threat intelligence, and vendor disclosures.

Step 2: Category Calibration

Raw dimension scores are calibrated against empirically-derived category risk profiles. Each of the 17 vulnerability categories has a documented typical score range based on observed incidents. The calibration formula blends analyst assessment with the empirical anchor:

```
Calibrated_Score = (Raw_Score x 0.55) + (Category_Anchor x 0.45)
```

Where Category_Anchor is the median composite score for that category across all observed incidents. This calibration ensures scoring consistency: our analysts assessing the same vulnerability type will converge toward the same score, anchored by empirical baselines.

Step 3: Composite Score

The composite score is the direct sum of calibrated dimension scores:

```
ASI-VSS = D1 + D2 + D3 + D4 + D5
```

No weights are applied. Each dimension contributes equally, reflecting that technical severity, AI amplification, propagation, enterprise impact, and exploitability are co-equal determinants of vulnerability severity.

Step 4: Classification

The composite score maps to the five-tier severity classification system defined in Chapter 3, determining the response timeline and insurance implications.

Step 5: CVSS Equivalent

For interoperability with CVSS-integrated vulnerability management platforms:

```
CVSS_Equivalent = ASI-VSS / 10    (range: 0.0 - 10.0)
```

PART III

Vulnerability Taxonomy

Chapter 5

17-Category Classification System

ASI-VSS classifies AI vulnerabilities into 17 empirically-derived categories. Each category is mapped to MITRE ATLAS technique identifiers, NIST AI RMF functions, and OWASP AIVSS risk categories where equivalent mappings exist. Two categories have no OWASP AIVSS equivalent, and a third—Shadow AI—extends significantly beyond its nearest OWASP mapping.

#	ASI-VSS Category	MITRE ATLAS	OWASP AIVSS	NIST AI RMF
1	Configuration Failures	AML.T0040	Agent Access Control Violation	GOVERN 1.2
2	Data Exfiltration	AML.T0024, AML.T0025	Agent Data Exfiltration	MAP 1.5
3	AI-Enabled Fraud	AML.T0048	Agentic AI Tool Misuse	MANAGE 2.4
4	Regulatory Violations	—	NO EQUIVALENT	GOVERN 1.1
5	Agent Compromise	AML.T0054	Agent Goal Manipulation	MANAGE 4.1
6	Code Vulnerabilities	AML.T0040	Agent Supply Chain & Dependency Attacks	MAP 3.4
7	Supply Chain Attacks	AML.T0010, AML.T0042	Agent Supply Chain & Dependency Attacks	MAP 3.1
8	Prompt Injection	AML.T0051, AML.T0043	Agent Goal & Instruction Manipulation	MEASURE 2.7
9	Adversarial Evasion	AML.T0015, AML.T0044	NO EQUIVALENT	MEASURE 2.6
10	MCP Protocol Exploits	AML.T0042	Agent Supply Chain & Dependency Attacks	MAP 3.4
11	Shadow AI	—	Agent Untraceability	GOVERN 1.5

#	ASI-VSS Category	MITRE ATLAS	OWASP AIVSS	NIST AI RMF
12	Model Theft & Extraction	AML.T0024, AML.T0044	Agent Data Exfiltration	MANAGE 2.2
13	Training Data Poisoning	AML.T0020	Agent Memory & Context Manipulation	MEASURE 2.5
14	RAG Poisoning	AML.T0020, AML.T0043	Agent Memory & Context Manipulation	MEASURE 2.7
15	Cascading Agent Failures	AML.T0054	Multi-Agent Orchestration Attacks	MANAGE 4.2
16	Memory & Context Poisoning	AML.T0020	Context Amnesia Exploitation	MEASURE 2.5
17	Model Denial of Service	AML.T0029	Agentic AI Tool Misuse	MANAGE 4.1

Table 5.1 — ASI-VSS 17-category taxonomy with framework mappings. Categories marked “NO EQUIVALENT” have no OWASP AIVSS coverage.

> TAXONOMY GAP ANALYSIS
>
> Two ASI-VSS categories — Regulatory Violations and Adversarial Evasion — have no OWASP AIVSS equivalent. Regulatory Violations encompasses EU AI Act and GDPR-triggered incidents (14 in our dataset, avg score 44.1) and Adversarial Evasion covers model-level attack techniques (6 incidents, avg 40.0). A third category, Shadow AI (8 incidents, avg 43.0), extends well beyond OWASP’s Agent Untraceability—covering unauthorized enterprise deployments, stolen AI account markets, credential exposure through AI tools, and organizational governance failures that transcend agent audit trails.

AIRS Domain Integration

Each ASI-VSS category maps to one or more ASI-AIRS (AI Insurance Readiness Score) domains, enabling carriers to assess which insurance coverage domains are implicated by specific vulnerability scores.

AIRS Domain	Code	VSS Categories Mapped
Model Integrity	MI	Training Data Poisoning, RAG Poisoning, Adversarial Evasion, Memory Poisoning, Model Theft
Operational Resilience	OR	Cascading Failures, Agent Compromise, MCP Exploits, Model DoS, Config Failures

AIRS Domain	Code	VSS Categories Mapped
Compliance & Legal	CL	Regulatory Violations, Shadow AI, Data Exfiltration, AI Fraud
Data Governance	DG	Data Exfiltration, Training Data Poisoning, Memory Poisoning
Third-Party Risk	TPR	Supply Chain Attacks, MCP Exploits, Code Vulnerabilities

Table 5.2 — AIRS domain integration mapping

PART IV

Empirical Foundation

Chapter 6

The 174-Incident Dataset

ASI-VSS is the first AI vulnerability scoring standard built on real-world operational data. Our intelligence team collected, verified, and scored 174 AI security incidents between January 2025 and March 2026, spanning all 17 vulnerability categories.

Dataset Overview

Metric	Value
Total incidents	174
Collection period	January 2025 – March 2026
Vulnerability categories	17
Score range	35–87
Mean score	55.3
Standard deviation	11.1
CVSS equivalent range	3.5–8.7

Table 6.1 — Dataset overview statistics

Severity Distribution

The severity distribution reflects the maturity of the AI security landscape: the majority of incidents (103, or 59%) fall in the ELEVATED tier, with a significant MODERATE tail (50 incidents). SEVERE incidents (19) represent high-impact events requiring rapid response, while 2 CRITICAL incidents demanded immediate action.

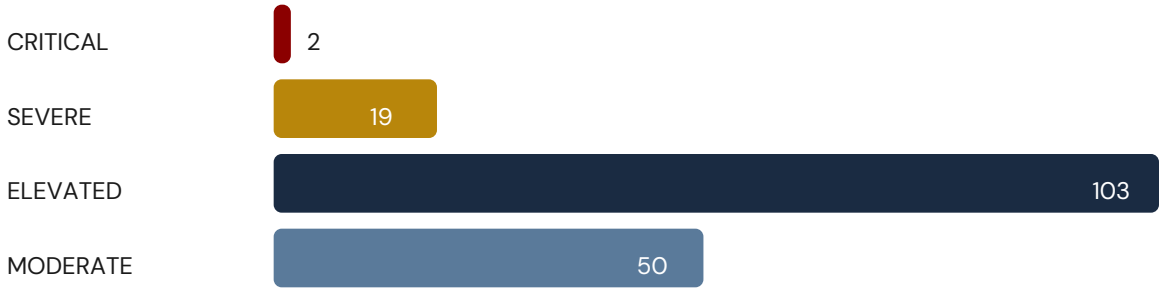

Figure 6.1 — Severity distribution across 174 incidents

Category Risk Profiles

Category	Avg Score	Count	Range	Trend
Agent Compromise	80.5	6	74–87	↑↑ Accelerating
Cascading Agent Failures	80.3	3	80–81	↑↑ Accelerating
Memory & Context Poisoning	75.0	1	75–75	↑ Emerging
MCP Protocol Exploits	69.9	8	62–85	↑↑ Accelerating
Supply Chain Attacks	68.9	14	62–78	↑ Emerging
RAG Poisoning	62.8	4	52–75	↑ Emerging
Training Data Poisoning	62.3	6	57–67	↑ Emerging
Data Exfiltration	57.8	4	51–63	→ Sustained
AI-Enabled Fraud	54.0	61	46–64	→ Sustained
Prompt Injection	53.2	10	46–60	↓ Plateauing
Model Theft & Extraction	50.8	6	43–61	→ Sustained
Configuration Failures	50.5	14	41–59	↓ Plateauing
Code Vulnerabilities	45.0	8	38–52	↓ Plateauing
Regulatory Violations	44.1	14	35–52	→ Sustained
Shadow AI	43.0	8	37–50	→ Sustained

Category	Avg Score	Count	Range	Trend
Adversarial Evasion	40.0	6	35–48	↓ Plateauing
Model Denial of Service	38.0	1	38–38	— Stable

Table 6.2 — Category risk profiles derived from 174 incidents, ordered by average score

Top 10 Highest-Scoring Incidents

VSS ID	Score	CVSS	Class.	Title
ASI-VSS-2026-0001	87	8.7	CRITICAL	CISA Warns of Active Exploitation of Critical Langflow Vulnera…
ASI-VSS-2026-0002	85	8.5	CRITICAL	n8n Workflow Automation Platform Remote Code Execution via Exp…
ASI-VSS-2026-0003	84	8.4	SEVERE	Moltbook AI Agent Network Breach — 1.5M Agents
ASI-VSS-2026-0004	83	8.3	SEVERE	Agent Credential Self-Generation
ASI-VSS-2025-0005	81	8.1	SEVERE	Multi-Agent Trust Boundary Violation
ASI-VSS-2025-0006	80	8.0	SEVERE	Cascading Agent Failure — 87% Downstream Corruption
ASI-VSS-2026-0007	80	8.0	SEVERE	OMNI-LEAK — Multi-Agent Data Collusion
ASI-VSS-2026-0008	78	7.8	SEVERE	OpenClaw Malicious Skill Trap: Hostile Skills Deploy Malware f…
ASI-VSS-2026-0009	78	7.8	SEVERE	Langflow Code Injection Vulnerability Enables Unauthenticated …
ASI-VSS-2026-0010	78	7.8	SEVERE	TeamPCP Compromises LiteLLM Supply Chain via CI/CD Pipeline, D…

Table 6.3 — Top 10 highest-scoring incidents in the ASI-VSS dataset

Dimension Analysis

Average dimension scores across all 174 incidents reveal where AI vulnerability severity concentrates:

Dimension	Average Score	Maximum	% of Maximum
D1: Technical Exploit Severity	11.4	20	57%
D2: AI Capability Amplification	6.4	20	32%

Dimension	Average Score	Maximum	% of Maximum
D3: Propagation & Persistence	10.6	20	53%
D4: Enterprise & Financial Impact	13.4	20	67%
D5: Exposure & Exploitability	13.4	20	67%

Table 6.4 — Average dimension scores across 174 incidents

DIMENSION INSIGHT

D4 (Enterprise & Financial Impact) and D5 (Exposure & Exploitability) lead at 13.4 and 13.4 respectively, reflecting that AI incidents consistently trigger significant organizational consequences. D2 (AI Capability Amplification) averages 6.4, indicating that while AI amplification is present, it varies widely by category — concentrated in agent compromise and cascading failure scenarios.

AIRS Domain Impact Analysis

Aggregating incident scores by AIRS domain reveals which insurance coverage areas face the highest concentration of AI vulnerability risk:

AIRS Domain	Code	Avg Score	Incidents	Range
Data Protection	DP	66.9	14	44–85
Third-Party Risk	TPR	59.2	30	38–84
Data Architecture	DA	59.0	55	43–87
Model Risk	MR	57.3	58	47–87
Data Governance	DG	56.1	34	38–81
Model Integrity	MI	55.3	79	35–87
Operational Resilience	OR	54.7	48	35–84
Regulatory Compliance	RC	54.2	37	43–63
Compliance & Legal	CL	48.5	45	35–64

Table 6.5 — AIRS domain impact analysis across 174 incidents

Chapter 7

Landscape & Structural Comparison

The AI vulnerability scoring landscape is evolving rapidly. This chapter provides a structural comparison between ASI-VSS v1.0 and OWASP AIVSS v0.8—two complementary frameworks working to define how the industry assesses AI-specific vulnerabilities.

Capability	ASI-VSS v1.0	OWASP AIVSS v0.8
Score range	0–100	0–10
AI amplification factors	5 continuous sub-factors (0–4 each)	10 ternary factors (0/0.5/1)
Propagation modeling	Dedicated D3 dimension (4 sub-factors)	Not addressed
Insurance loss mapping	D4 sub-factor with coverage triggers	Not addressed
Empirical data	174 analyst-verified incidents	Expert surveys + community input
Vulnerability categories	17	10
Category calibration	Empirical anchoring (55/45 blend)	Not addressed
MITRE ATLAS mapping	Direct technique IDs	Indirect
NIST AI RMF mapping	Direct function codes	Indirect via MAESTRO
Severity tiers	5 tiers with insurance implications	4 tiers (standard CVSS)
CVSS interoperability	VSS/10 = CVSS equivalent	Native 0–10 (additive)
Release status	v1.0 — April 15, 2026	v0.8 pre-release; v1.0 review opens April 16

Table 7.1 — Structural comparison: ASI-VSS v1.0 vs. OWASP AIVSS v0.8

Where ASI-VSS Extends the Conversation

Beyond granularity and empirical grounding, ASI-VSS provides structural capabilities that extend beyond OWASP AIVSS:

Propagation modeling. ASI-VSS Dimension 3 scores persistence duration, propagation velocity, cross-system reach, and detection difficulty—four sub-factors that characterize how AI vulnerabilities spread through interconnected systems—an area where OWASP AIVSS may expand in future versions.

Insurance loss mapping. The D4 Insurance Loss Category sub-factor maps vulnerability severity directly to cyber insurance coverage triggers (Business Interruption, Cyber Liability, D&O, Data Breach Notification). This enables carriers to assess coverage implications from VSS scores without secondary analysis.

Empirical calibration. ASI-VSS's category calibration mechanism anchors analyst assessments to baselines derived from 174 observed incidents, complementing the expert consensus approach that OWASP AIVSS has built from its community engagement.

Seven additional categories. ASI-VSS covers 17 vulnerability categories versus OWASP's 10, including Regulatory Violations, Adversarial Evasion, Shadow AI, MCP Protocol Exploits, RAG Poisoning, Memory & Context Poisoning, and Model Denial of Service—categories that represent real and growing threat vectors in production AI systems.

PART V

Adoption & Integration

Chapter 8

Implementation Guidance

ASI-VSS is designed for immediate operational integration across three primary stakeholder groups: security practitioners, insurance carriers, and regulatory bodies.

For Practitioners

Security teams can integrate ASI-VSS scores directly into vulnerability management workflows. The five-dimension architecture allows triage based on specific risk dimensions—an incident scoring high on D3 (Propagation) demands different containment strategies than one scoring high on D4 (Enterprise Impact). The CVSS equivalent (VSS/10) enables seamless integration with existing vulnerability management platforms that accept CVSS inputs.

Our analysts recommend the following integration approach:

- Ingest VSS scores alongside CVSS scores in vulnerability management platforms
- Use dimensional breakdowns to route incidents to appropriate response teams (D2-heavy scores to AI/ML teams, D4-heavy scores to legal/compliance)
- Apply severity tiers to SLA definitions—CRITICAL and SEVERE incidents map directly to response timelines
- Track category trends to anticipate emerging threat concentrations

For Carriers

Insurance carriers can leverage ASI-VSS to enhance underwriting precision for AI-related risks. The D4 Insurance Loss Category sub-factor maps directly to standard coverage triggers, enabling carriers to assess portfolio exposure by vulnerability category.

- Underwriting integration: Use category risk profiles (Table 6.2) to calibrate AI risk premiums by vulnerability type
- Claims assessment: Reference D4 scores and insurance loss categories when evaluating AI-related claims
- Portfolio analysis: Aggregate AIRS domain scores (Table 6.5) to identify coverage concentration risks
- Renewal decisions: Track VSS severity trends to inform renewal pricing adjustments

For Regulators

ASI-VSS aligns with emerging regulatory frameworks for AI risk management:

• EU AI Act: VSS severity classifications map to the Act's risk tiers for high-risk AI systems. D4 Regulatory Exposure scoring tracks compliance triggers explicitly
• NIST AI RMF: All 17 VSS categories are mapped to NIST AI RMF functions (GOVERN, MAP, MEASURE, MANAGE), enabling alignment with the federal framework
• SEC AI Disclosure: VSS scores provide quantifiable metrics for material AI risk disclosure requirements

ASI-AIRS Integration

ASI-VSS feeds directly into our AI Insurance Readiness Score (AIRS) platform. VSS scores contribute to AIRS domain calculations—Model Integrity, Operational Resilience, Compliance & Legal, Data Governance, and Third-Party Risk—providing carriers with a unified view of an organization's AI security posture alongside vulnerability-specific severity assessments.

APPENDIX

Reference Materials

Full Methodology Citation

```
AI Security Intelligence. (2026). ASI-VSS: Vulnerability Scoring Standard v1.0.
AI Security Intelligence LLC.
https://aisecurityintelligence.com/vss
```

Glossary of Terms

Term	Definition
ASI-VSS	AI Security Intelligence Vulnerability Scoring Standard
AIRS	AI Insurance Readiness Score — ASI's platform for AI risk assessment
CVSS	Common Vulnerability Scoring System (FIRST.org, currently v4.0)
AIVSS	AI Vulnerability Scoring System (OWASP, currently v0.8 pre-release)
MITRE ATLAS	Adversarial Threat Landscape for AI Systems — MITRE's AI attack framework
NIST AI RMF	NIST Artificial Intelligence Risk Management Framework
D1–D5	The five scoring dimensions of ASI-VSS (0–20 each)
Category Anchor	Median composite score for a vulnerability category across observed incidents
MCP	Model Context Protocol — standardized protocol for AI tool integration
RAG	Retrieval-Augmented Generation — architecture pattern combining LLMs with external data
CIA Triad	Confidentiality, Integrity, Availability — foundational security model
D&O;	Directors and Officers liability insurance

Table A.1 — Glossary of terms

Contact

AI Security Intelligence LLC
aisecurityintelligence.com/vss

ASI-VSS is maintained by AI Security Intelligence LLC. The standard is freely available for adoption. Scored vulnerability data and continuous monitoring are available through the ASI intelligence platform.