ASI-VSS

Vulnerability Scoring Standard v1.0

Practitioner Brief

Why Traditional Scoring Falls Short

The security industry has relied on the Common Vulnerability Scoring System (CVSS) for more than two decades. CVSS was designed for deterministic software: a buffer overflow either exists or it does not; a remote code execution either requires authentication or it does not. The scoring logic assumes static, bounded systems where attack paths are fixed and outcomes predictable. AI systems violate every one of these assumptions.

Modern AI deployments introduce vulnerabilities that are fundamentally non-deterministic. A prompt injection attack may succeed on one inference pass and fail on the next, depending on context window state, model temperature, or retrieved document composition. Retrieval-Augmented Generation (RAG) poisoning corrupts knowledge bases serving thousands of downstream queries. Model Context Protocol (MCP) exploits traverse tool chains that did not exist when traditional scoring rubrics were written. Autonomous agent architectures can propagate compromise at machine speed, cascading failures through multi-agent pipelines in ways that no CVSS base metric captures.

The gap is not merely theoretical. Our analysts observe that the vulnerability categories generating the most severe real-world outcomes — agent compromise, cascading agent failures, and memory and context poisoning — produce CVSS scores of 7–9 when force-mapped, obscuring the critical operational distinction between a targeted credential theft and an autonomous AI system actively exfiltrating data across an enterprise. The scoring delta matters enormously for response prioritization, insurance coverage triggers, and regulatory disclosure.

OWASP AIVSS v0.8 took an important step forward with 10 agentic factors designed for AI system risk. ASI-VSS builds on this foundation with five continuous scoring dimensions, 20 granular sub-factors, and empirical calibration against 174 analyst-verified AI security incidents collected January 2025 through March 2026 — a scoring standard that speaks the language of AI infrastructure risk.

> KEY FINDING
> The highest-scoring vulnerability categories in our incident dataset are 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 a scoring architecture that reflects autonomous propagation, non-deterministic behavior, and multi-agent cascading — dimensions absent from CVSS and only partially addressed in prior AI-specific frameworks.

Five-Dimension Scoring Architecture

ASI-VSS decomposes AI vulnerability severity into five independent dimensions, each scored 0–20 by trained analysts using structured sub-factor rubrics. The composite score reflects both technical severity and AI-specific amplification factors that deterministic frameworks cannot capture.

```
VSS = D1 + D2 + D3 + D4 + D5        (scale: 0 – 100)
```

Dimension	Range	Sub-F.	Description
D1 Technical Exploit Severity	0–20	4	CIA impact across confidentiality, integrity, and availability; attack complexity and privilege requirements
D2 AI Capability Amplification	0–20	5	Autonomy level, non-deterministic behavior, tool-use reach, learning and adaptation, multi-agent exposure
D3 Propagation & Persistence	0–20	4	Duration of exposure, propagation velocity, organizational reach, detection difficulty and dwell time
D4 Enterprise & Financial Impact	0–20	4	Quantified financial loss, regulatory and compliance exposure, reputational damage, insurance coverage mapping
D5 Exposure & Exploitability	0–20	4	Exploit maturity and availability, vendor defense posture, attack surface breadth, analyst confidence rating

CVSS Interoperability

CVSS Equivalent = VSS ÷ 10 — ASI-VSS maps natively to the CVSS 10-point scale for integration with existing vulnerability management workflows, SIEM platforms, and risk registers.

LOW (0–19) · MODERATE (20–39) · ELEVATED (40–59) · HIGH (60–79) · CRITICAL (80–100)

Grounded in 174 Real-World Incidents

ASI-VSS is calibrated against 174 analyst-verified AI security incidents collected by our intelligence team from January 2025 through March 2026. Each incident was independently scored by our analysts using the full ASI-VSS sub-factor rubrics, with disagreements resolved through structured adjudication. The dataset spans 17 vulnerability categories mapped to MITRE ATLAS technique identifiers, NIST AI RMF functions, and OWASP risk classifications.

DATASET DISTRIBUTION

Score Range	Mean Score	Std Dev	Critical	Severe	Elevated	Moderate
35 – 87	55.3	11.1	2	19	103	50

Top Vulnerability Categories by Average Score

Vulnerability Category	Avg Score	Incidents	Score Range
Agent Compromise	80.5	6	74 – 87
Cascading Agent Failures	80.3	3	80 – 81
Memory & Context Poisoning	75.0	1	75 – 75
MCP Protocol Exploits	69.9	8	62 – 85
Supply Chain Attacks	68.9	14	62 – 78
RAG Poisoning	62.8	4	52 – 75
Training Data Poisoning	62.3	6	57 – 67
Data Exfiltration	57.8	4	51 – 63

The full 17-category taxonomy maps each incident to MITRE ATLAS technique identifiers, NIST AI RMF governance functions, and OWASP risk classification codes — enabling practitioners to cross-reference ASI-VSS scores with existing threat intelligence and compliance frameworks.

Implementation Guide

ASI-VSS is designed for operational deployment by security practitioners, GRC teams, and AI red teamers. The guidance below covers the primary use cases our analysts support across enterprise, insurance, and regulatory contexts.

Vulnerability Assessment

Score each dimension using the sub-factor rubrics in the full ASI-VSS v1.0 Specification. Each sub-factor is rated on a normalized scale; dimension scores are summed to produce the composite VSS. Dual-analyst scoring with structured adjudication is recommended for any vulnerability above VSS 60.

Severity Classification

Apply the five-tier system — CRITICAL (80–100), SEVERE (60–79), ELEVATED (40–59), MODERATE (20–39), LOW (0–19) — to drive response timelines and escalation protocols. CRITICAL and SEVERE findings warrant immediate notification to security leadership and regulatory disclosure review where applicable.

Framework Interoperability

ASI-VSS maps natively to CVSS (composite ÷ 10), MITRE ATLAS technique identifiers for AI-specific tactics and techniques, and NIST AI RMF governance functions (Map, Measure, Manage, Govern). Existing vulnerability management platforms, SIEMs, and risk registers can ingest ASI-VSS scores alongside traditional CVSS data without workflow disruption.

Insurance Integration

D4 sub-factors — financial loss quantification, regulatory exposure, reputational damage, and coverage-trigger mapping — align directly with cyber insurance underwriting criteria. VSS scores feed into the AI Insurance Readiness Score (AIRS) for underwriting assessment. Carriers and brokers using AIRS can request scored incident benchmarking from our intelligence team.

Continuous Monitoring

AI threat landscapes evolve faster than traditional software ecosystems. Our analysts recommend rescoring high-severity findings quarterly, or immediately following material changes to model version, tool integrations, or deployment architecture. Category calibrations update with each refresh of the ASI-VSS incident database, keeping scores anchored to observed real-world impact.

Get Started

ASI-VSS v1.0 is available to security practitioners, insurance carriers, and regulatory bodies under licensing arrangements suited to organizational scale and use case. The following resources are available from our intelligence team.

- 1. Full ASI-VSS v1.0 Specification
  Complete scoring rubrics, sub-factor definitions, calibration methodology, and analyst guidance. The Specification is the authoritative reference for all scoring decisions and framework integrations.
- 2. ASI-VSS Scored Incident Database
  All 174 analyst-verified AI security incidents, fully scored and annotated with MITRE ATLAS technique IDs, NIST AI RMF function mappings, and OWASP risk classifications. Available for benchmarking, calibration validation, and comparative analysis.
- 3. Contact Our Intelligence Team
  Licensing inquiries, platform integration support, custom vulnerability assessments, and AIRS underwriting consultations. Our analysts are available for engagements ranging from single-assessment support to ongoing embedded intelligence partnerships.

> Contact
> info@aisecurityintelligence.com · aisecurityintelligence.com

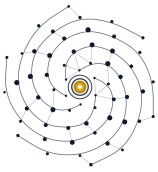

About AI Security Intelligence LLC

AI Security Intelligence LLC provides the intelligence infrastructure for quantifying AI security risk. Our Vulnerability Scoring Standard, Threat Severity Score, and AI Insurance Readiness Score form an integrated platform serving security practitioners, insurance carriers, and regulators worldwide. Our analysts combine deep expertise in AI systems security, enterprise risk management, and actuarial modeling to deliver intelligence that is both technically rigorous and operationally actionable.