

AIMRS™

Framework Overview

A practical management architecture for intelligent systems

PUBLICATION STATUS Official public release. This publication provides high-level AIMRS™ guidance. Controlled scoring keys, control applicability logic, assessment instruments, assurance procedures and practitioner materials are not included.

Originated and stewarded by CyberVirtues

Public Release 1.0 | August 22, 2026

Document ID: AIMRS-PUB-FWK-1.0

© 2026 Cyber Virtues LLC. All rights reserved.

1. What AIMRS™ Is

AIMRS™ - Artificial Intelligence Management, Risk & Security - is a technology-neutral framework for directing, governing, securing, operating and measuring AI capabilities that perform or influence work.

CORE IDEA AI must be managed - not merely implemented.

The framework applies three integrated lenses - Management, Risk and Security - across seven stable domains. It is designed for AI applications, human-AI teams, agents, automation, robotics and other intelligent systems throughout their lifecycle.

2. Foundational Principles

Principle	Management commitment
Human accountability	A named person or governing body remains accountable for every material AI capability.
Explicit authority	AI acts only within approved, defined and observable boundaries.
Proportional management	Requirements scale with authority, impact, sensitivity, reach, reversibility and uncertainty.
Lifecycle control	AI is controlled continuously from proposal through retirement.
Secure by design	Identity, data, models, tools, infrastructure and actions are protected as one operating whole.
Evidence over assertion	Trust and conformance are supported by records, testing and monitoring.
Operational resilience	Failure, attack, drift, outage and unexpected behavior are anticipated and recoverable.
Value with responsibility	AI produces legitimate value without unacceptable or unmanaged impact.
Interoperability	Implementation maps to established standards and avoids unnecessary duplication.
Continuous evolution	Management adapts as capability, context, threats and obligations change.

3. Seven Management Domains

Domain	Objective
Strategy	Direct AI toward mission, value and risk objectives.
Workforce	Design effective and accountable human-AI work systems.
Governance	Establish authority, ownership, policy and oversight.
Security	Protect AI capabilities and their operating environment.
Operations	Operate AI reliably as part of daily business.
Performance	Measure outcomes, behavior, cost, quality and impact.
Lifecycle	Control entry, change, suspension and exit.

Each domain is considered through six cross-cutting lenses: people, process, technology, data, third parties and external impact.

4. Implementation Lifecycle

View	Executive question
Align	Why are we using AI, what value is expected and what boundaries apply?
Govern	Who owns decisions, authority, policy, risk and exceptions?
Assure	What evidence shows that design and operation meet expectations?
Operate	How is AI monitored, supported, changed, interrupted and improved?

ARCHITECTURE NOTE Align, Govern, Assure and Operate are a simplified implementation lifecycle view. They do not replace the seven AIMRS™ management domains.

5. Maturity Model

Level	Name	Organizational condition
0	Baseline	Scope is defined, the AI estate is visible enough to identify critical gaps, and a prioritized improvement plan exists.
1	Managed	Ownership, classification, controls and operating practices are defined, assigned and operating across the required scope.
2	Assured	Operating effectiveness is measured and independently challenged in proportion to risk.
3	Adaptive	Trends, learning and evidence drive anticipatory and continual improvement.

- Maturity measures organizational capability, not model sophistication.
- A high average does not compensate for a critical domain at an unacceptable level.
- The appropriate target depends on context; Adaptive maturity is not required everywhere.
- Assessment should drive prioritized improvement rather than ceremonial scoring.

6. Who Can Use AIMRS™

AIMRS™ is designed for organizations of any size. Small organizations can establish a practical baseline using a narrow scope and a few essential actions. Larger organizations can add portfolio, business-unit, jurisdictional and technology overlays without changing the core architecture.

Organization profile	Recommended starting point
Small: 1-99	One accountable owner, practical inventory, material-use review and priority actions.
Mid-size: 100-999	Cross-functional ownership, portfolio visibility, lifecycle gates and regular review.
Large: 1,000-4,999	Federated ownership, formal assurance, specialized overlays and enterprise reporting.
Enterprise: 5,000+	Multi-entity governance, shared minimums, local accountability, independent challenge and aggregated evidence.

Publication Notice

AIMRS™ (Artificial Intelligence Management, Risk & Security) is developed and stewarded by Cyber Virtues LLC. AIMRS™ and associated names, marks, framework structures and publications are proprietary intellectual property of Cyber Virtues LLC.

Permitted public use

- Download, read, cite and link to the unmodified publication for internal education and evaluation.
- Reference AIMRS™ accurately with attribution to Cyber Virtues LLC and the applicable release number.
- Use the publication to inform internal planning without representing independent certification or authorization.

Restricted use

- Do not modify, rebrand, resell, sublicense or publish derivative standards under the AIMRS™ name without written permission.
- Do not claim AIMRS™ certification, assessor status, conformance or endorsement unless formally authorized through an operational AIMRS™ program.
- Do not reproduce controlled scoring keys, assessment instruments, control logic or practitioner-only materials.

VERSION CONTROL Verify that this is the current public release at cybervirtues.com/ai-framework. References should identify the publication title and Public Release 1.0.