DATACENDIA, INC.

NIST AI Risk Management Framework Alignment

Voluntary Alignment Statement

Organization: Datacendia, Inc.
Document Version: 1.0
Date: February 6, 2026
Framework Referenced: NIST AI RMF 1.0 (January 2023)
Classification: Public

Purpose

This document maps Datacendia's AI decision governance platform to the four core functions of the NIST AI Risk Management Framework. This is a voluntary alignment statement, not a certification.

Framework Mapping

GOVERN — Establish and maintain AI governance

GOVERN Subcategory	Datacendia Implementation
Policies and procedures	Decision accountability enforced at architecture level; governance policies codified as executable rules, not documents
Roles and responsibilities	Role-based authority model with explicit override permissions; every decision records who authorized it and under what authority
Risk culture	Multi-agent deliberation designed to surface disagreement before execution; dissent is a first-class artifact, not an exception
Legal and regulatory awareness	Jurisdiction-aware processing; cross-jurisdiction compliance engine covering 17 regulatory regimes; regulatory framework mapping per decision

MAP — Contextualize AI risks

MAP Subcategory	Datacendia Implementation
System context	Intent, constraints, and jurisdiction captured pre-decision; organizational context scoped per tenant
Stakeholder identification	Agent perspectives represent distinct stakeholder viewpoints (legal, financial, ethical, operational, adversarial)
Risk identification	Pre-execution risk surfacing via multi-perspective deliberation; adversarial red team mode available for stress-testing
Benefit-risk assessment	Trust Delta computation quantifies governance overhead vs. decision quality improvement

MEASURE — Analyze and assess AI risks

MEASURE Subcategory	Datacendia Implementation
Risk metrics	Confidence scoring, dissent severity tracking, risk heat mapping across decision categories
Bias assessment	Multi-perspective agent architecture designed to surface bias through structured disagreement; minority harm analysis available
Performance monitoring	Decision outcome tracking; agent contribution analysis; governance effectiveness metrics
Transparency artifacts	Full deliberation transcripts with evidence citations; reasoning chains preserved and replayable

MANAGE — Prioritize and act on AI risks

MANAGE Subcategory	Datacendia Implementation
Risk treatment	Escalation paths enforce human review for high-risk decisions; policy gates prevent execution without required approvals
Incident response	Override documentation with justification requirements; tamper-evident audit ledger for post-incident review
Communication	Regulatory receipt generation for external reporting; structured decision packets for audit consumption
Continuous improvement	Decision replay capability for post-hoc analysis; governance metrics tracked for trend identification

Key Architectural Principle

Datacendia does not replace domain-specific risk controls.

Datacendia records and proves them.

The platform produces regulator-grade evidence that governance processes were followed, disagreements were captured, and human oversight was exercised.

Verification

All claims in this document can be independently verified using:

- Cryptographic Decision DNA artifacts (Merkle root integrity, hash chain verification)
- Open verification tooling (available at datacendia.com/trust)
- Full deliberation transcripts with timestamps and agent attribution

Planned Actions

Action	Target Date
Formal NIST AI RMF self-assessment using AI RMF Playbook	Q2 2026
Third-party NIST alignment review	Q4 2026
Publication of NIST AI RMF Playbook responses	Q2 2026

Stuart Rainey
Founder & CEO
Datacendia, Inc.
February 6, 2026

DOCUMENT INTEGRITY
SHA-256: 739da41d7253e28ffb6cda69e05db186b79cac32d67de366c408d9194aba89b5
This hash covers the full text content of this document. Verify at datacendia.com/trust