

# Responsible AI Statement

Daedalus AI Holdings LLC (“Daedalus”) develops artificial intelligence-enabled systems designed to support enterprise software development, automation, configuration management, and technical governance workflows. Our systems are intended for use in professional and institutional environments where reliability, oversight, and security are essential.

Artificial intelligence technologies introduce meaningful benefits alongside operational and governance risks. Our approach to Responsible AI reflects a structured, risk-aware philosophy grounded in engineering discipline, defined system boundaries, and explicit allocation of responsibility between technology provider and customer.

This statement describes the principles that guide the design, deployment, and oversight of AI-enabled features within our Services.

## 1. Purpose and Intended Scope of AI Systems

Daedalus’ AI systems are designed to assist with technical workflows such as software development, code refactoring, documentation generation, configuration design, system diagnostics, and automation of structured engineering tasks.

The intended purpose of our AI systems is technical augmentation, not automated adjudication. Our systems are not designed or marketed to:

- A.** Evaluate, score, rank, or profile natural persons
- B.** Make employment, credit, housing, insurance, or similar consequential decisions
- C.** Conduct biometric identification or surveillance
- D.** Operate as autonomous decision-making systems without human review

## 2. Human Oversight and Professional Judgment

We design our AI-enabled features to operate as assistive tools for qualified professionals. Outputs generated by our systems require independent review, validation, testing, and approval prior to operational deployment.

AI-generated content may include code, documentation, configurations, recommendations, or analytical summaries. Such outputs are probabilistic in nature and may contain inaccuracies, logical errors, incomplete reasoning, or security vulnerabilities.

Daedalus does not represent that AI-generated outputs are error-free, secure, complete, or suitable for unreviewed production deployment. Responsible AI use requires structured human governance. Customers remain responsible for:

- A.** Validating technical correctness
- B.** Conducting quality assurance testing
- C.** Performing security review
- D.** Ensuring compliance with applicable regulatory requirements
- E.** Determining suitability for production environments

### **3. AI Transparency and System Characteristics**

We believe transparency requires clarity regarding system purpose, boundaries, and limitations.

#### **3.1 Model Characteristics**

Our Services may incorporate pretrained language models and proprietary orchestration layers that coordinate system interactions. Outputs are generated using probabilistic methods and may vary across sessions or inputs.

Transparency does not require disclosure of proprietary architecture, model weights, training datasets, internal prompt structures, or security-sensitive implementation details.

#### **3.2 Output Variability and Limitations**

AI outputs are influenced by model architecture, prompt structure, system context, and probabilistic inference. These characteristics are inherent to large language model systems. Outputs may:

- A.** Differ across similar prompts
- B.** Reflect incomplete reasoning chains
- C.** Include outdated or incorrect information
- D.** Omit context not provided in the input

#### **3.3 No Claim of Deterministic Explainability**

Daedalus does not represent that all AI outputs are deterministically explainable at the token-by-token level. Where appropriate, system behavior can be contextualized at a functional level, but full algorithmic transparency is neither technically feasible nor operationally appropriate in many model architectures.

#### **3.4 Defined System Boundaries**

AI-generated outputs are produced within defined service boundaries. The system does not independently access external enterprise systems unless explicitly configured by the customer within the Service environment.

### **4. Data Boundaries and Model Training Practices**

Daedalus does not use customer-submitted enterprise data to train generalized artificial intelligence models. Customer inputs and enterprise data remain within defined service boundaries and are not incorporated into shared model training datasets.

We design logical separation mechanisms to prevent cross-customer data exposure within hosted environments. This commitment reflects our enterprise trust posture and our focus on institutional deployments.

### **5. Privacy and Data Governance Principles**

Where Daedalus processes personal information in connection with its Services, we do so within clearly defined roles and in accordance with applicable data protection laws.

We do not sell personal information, and we do not use enterprise customer data for advertising purposes.

We encourage customers to implement appropriate data classification, minimization, and governance controls prior to submitting data into AI-enabled workflows.

## 6. Security-Oriented AI Design

Responsible AI requires secure AI. We integrate security considerations into system architecture and operations, including:

- A. Encryption of data in transit and at rest
- B. Role-based access controls
- C. Least privilege access principles
- D. Monitoring and logging of administrative activity
- E. Structured change management
- F. Vulnerability management processes
- G. Incident response procedures

Security controls are designed to protect environments under our operational control. No system can eliminate all risk, but we prioritize minimizing operational and data security exposure.

## 7. Governance and Shared Responsibility

Effective AI governance requires clarity regarding responsibility allocation. Responsible AI is not achieved solely through technical safeguards; it requires institutional governance discipline.

### 7.1 Daedalus is responsible for:

- A. Designing and operating AI systems within defined service boundaries
- B. Implementing security controls within environments under our control
- C. Maintaining structured operational oversight processes

### 7.2 Customers are responsible for:

- A. Determining regulatory applicability
- B. Classifying and validating their data
- C. Reviewing AI-generated outputs
- D. Implementing internal compliance controls
- E. Ensuring proper human oversight

## 8. Fairness and Appropriate Use

Daedalus' AI systems are designed for technical and software-related workflows. They are not intended to evaluate individuals or to produce legally or socially consequential determinations. Where AI outputs may reflect limitations inherent in training data or model architecture, professional review remains essential. We do not represent that our systems eliminate bias or guarantee neutrality in all outputs.

## 9. Prohibited and High-Risk Uses

We do not support the use of our systems for unlawful, harmful, abusive, or security-compromising activities. This includes attempts to:

- A. Circumvent safeguards
- B. Extract proprietary system components
- C. Reverse engineer model weights
- D. Use the system for malicious or deceptive purposes

Organizations deploying AI-assisted workflows should conduct internal risk assessments appropriate to their use case.

## 10. Continuous Monitoring and Improvement

Artificial intelligence technologies and regulatory expectations evolve over time. We periodically review and refine our internal governance, security, and operational practices to reflect evolving industry standards and operational experience. Updates to this Responsible AI & Transparency Statement may occur to reflect such refinements.

## 11. Regulatory Context

Daedalus designs its AI systems for use in enterprise environments, including organizations subject to financial, healthcare, or governmental regulation. However, customers remain responsible for determining how our Services align with their specific regulatory obligations. Nothing in this statement constitutes a regulatory certification or authorization.

## 12. Ongoing Dialogue

We recognize that responsible AI governance is an evolving field. We welcome engagement from customers, regulators, and industry stakeholders regarding best practices in enterprise AI deployment. Questions regarding this Responsible AI & Transparency Statement may be directed to:

Daedalus AI Holdings LLC  
1919 West Greenleaf  
Chicago, Illinois 60626  
legal@trydaedalus.ai

**Last Updated:** March 2, 2026