

AI Governance & Responsible AI Overview

Version: 2.0

Last Updated: 1st September 2026

Classification: Public

AI Governance & Responsible AI Overview

Artificial intelligence offers real opportunities to improve efficiency, decision-making, and business innovation — and introduces real risks relating to governance, transparency, privacy, security, fairness, accountability, and regulatory compliance. This document describes Concord's approach to the responsible design, development, deployment, and operation of AI-enabled capabilities within the platform.

Our AI Philosophy

AI should augment human expertise, not replace it. Concord designs AI capabilities to support responsible business outcomes rather than autonomous decision-making.

Responsible AI Principles

Human oversight — AI-generated recommendations support human decision-makers. Customers remain responsible for reviewing AI outputs before relying on them for material legal, regulatory, financial, employment, healthcare, insurance, underwriting, or other consequential decisions.

Transparency — users should know when AI is being used. AI-generated content is always visibly labeled and never blended silently with human-authored content; drafted AI suggestions require a human action before they're posted or applied.

Accountability — organizations, not AI systems, remain accountable for decisions. Concord provides governance capabilities to assign ownership, document approvals, maintain audit trails, and demonstrate accountability across the AI lifecycle.

Privacy — AI capabilities are designed to operate consistently with our privacy commitments and contractual obligations relating to customer data.

Security — AI systems are developed and operated using the same security practices described in the [Security Overview](#), applied throughout the AI lifecycle.

Continuous improvement — we periodically review and enhance our AI governance practices as technology, customer needs, and regulatory expectations evolve.

AI Governance Framework

Concord supports governance across the AI lifecycle: strategy, AI inventory management, risk identification and assessment, policy management, human approval workflows, control implementation, compliance evidence, and ongoing monitoring — as an ongoing operational discipline rather than a one-time assessment.

Human-in-the-Loop

Organizations may configure workflows requiring human review before significant actions complete — policy approvals, AI risk assessments, governance decisions, compliance reviews, exception approvals, and model deployment approvals. The appropriate level of oversight depends on each customer's governance requirements and risk appetite.

AI Prompt and Output Handling

Customers retain ownership of prompts they submit and outputs generated for them, subject to the applicable Master Services Agreement. Concord processes prompts and outputs only as reasonably necessary to provide the Services, generate requested responses, maintain platform security, investigate operational issues, improve service quality using de-identified or aggregated information, and comply with legal obligations. Concord does not knowingly use identifiable customer prompts or customer data to train publicly available foundation AI models without the customer's prior authorization.

AI Risk Management

Organizations should evaluate AI risk in the context of their own use cases and requirements, considering accuracy, explainability, bias, fairness, privacy, security, reliability, human oversight, regulatory compliance, and business impact.

AI Model Management

Concord's AI-enabled capabilities are built on a multi-provider abstraction layer rather than a single hardcoded model — primarily Google Gemini, the same cloud ecosystem as hosting rather than a separate subprocessor relationship, with additional providers available for resilience. Because AI technologies continue to evolve rapidly, the specific underlying models may change over time; when they do, Concord aims to preserve core functionality, minimize customer disruption, and maintain contractual commitments.

AI Explainability

Where appropriate, Concord supports explainability through workflow transparency, governance records, audit trails, and decision-support information. The level of explainability available depends on the underlying AI technology and the nature of the task.

AI Lifecycle Governance

Concord supports governance activities across the full AI lifecycle: strategy, use-case registration, risk identification and assessment, policy review, human approval, deployment, operational monitoring, periodic reassessment, and retirement. Maintaining governance records throughout the lifecycle helps organizations demonstrate accountability and regulatory readiness.

Alignment with Industry Frameworks

Concord designs its AI governance capabilities with reference to the NIST AI Risk Management Framework, ISO/IEC 42001, ISO/IEC 23894, the OECD AI Principles, and

emerging AI regulation. References to these frameworks indicate alignment objectives, not certification, unless expressly stated.

Customer Responsibilities

Customers remain responsible for defining acceptable AI use cases, reviewing AI-generated outputs, establishing their own governance policies, managing user access, complying with applicable laws, providing appropriate human oversight, and determining whether AI outputs are suitable for their intended purpose.

Additional Information

Available on request, subject to confidentiality obligations: AI Services Terms & Responsible AI Policy, AI governance architecture information, [Security Overview](#), [Privacy & Data Protection Overview](#), and [Compliance Mapping](#).
