Artificial Intelligence Governance Policy

CredX Tech Inc

September 2026

Contents

Table 1: Control satisfaction

Standard	Controls Satisfied
TSC	CC1.1, CC3.2, CC5.1, CC5.2, CC6.1, CC8.1, CC9.2, P4.1

Table 2: Document history

Date	Comment
Sep 3 2026	Added automated credit decisioning governance

1 Purpose and Scope

a. The purpose of this policy is to establish requirements for the responsible, secure, and lawful use of artificial intelligence (AI) at CredX, whether embedded in products delivered to customers or used internally to support business operations.

b. This policy applies to all AI systems and services used by or on behalf of CredX, including generative AI services, machine learning models, AI-assisted development tools, and AI features embedded within third-party software.

c. This policy applies with particular force to **CredX's own machine learning models used in automated credit decisioning**. Those models determine whether an individual is offered credit, and are governed by the Automated Credit Decisioning Models section below in addition to every other requirement in this policy.

d. This policy applies to all employees, contractors, and vendors who develop, integrate, operate, or use AI systems in the course of CredX business.

2 Background

a. AI systems introduce risks that are not fully addressed by conventional application controls. Information submitted to an AI service leaves CredX's control boundary and may be retained or used by the provider. AI systems can produce output that is fluent but factually incorrect, and that output may be relied upon in downstream business decisions. Where AI processes customer or personal data, additional contractual and regulatory obligations apply.

b. CredX's use of AI is deliberately bounded. This policy defines those boundaries, the approval required to extend them, and the controls that apply in each case.

3 Definitions

a. **AI system**: Any system or service that uses machine learning or generative models to produce output, including large language models, whether operated by CredX or accessed as a third-party service.

b. **Inference**: Submitting input to a trained model to obtain output. Inference does not modify the model.

c. **Training or fine-tuning**: Using data to create or modify a model's parameters, such that the data influences future output.

d. **AI provider**: A third party operating an AI system accessed by CredX, including through an API or an embedded product feature.

e. **AI-assisted output**: Any content, code, data attribute, or recommendation generated in whole or in part by an AI system.

4 References

a. Data Classification Policy

b. Vendor Management Policy

c. Privacy Management Policy

d. Software Development Policy

e. System Change Policy

f. Processing Integrity Policy

5 Policy

5.1 AI Inventory and Approval

a. All AI systems and services used for CredX business must be approved by the CTO prior to use and recorded in the Asset Register in accordance with the Asset Management Policy, including the provider, purpose, data classification permitted, and business owner.

b. Employees and contractors must not use unapproved AI systems to process any CredX or customer information. Personal or trial accounts with AI services must not be used for CredX business.

c. The AI inventory must be reviewed at least annually, and whenever a new AI capability is introduced or an existing provider materially changes its terms.

5.2 Permitted Data by Classification

a. The following restrictions govern what information may be submitted to an AI system, in accordance with the classification levels defined in the Data Classification Policy:

   i. **Public**: Permitted.

   ii. **Internal Use**: Permitted with an approved AI system.

   iii. **Restricted**: Permitted only with an approved AI system whose provider terms have been verified under this policy, and only where necessary for a defined business purpose.

iv. **Confidential**: Prohibited unless expressly approved by the CTO, supported by verified provider terms, and permitted under the applicable customer agreement.

b. Personal information must not be submitted to any AI system unless a lawful basis exists under the Privacy Management Policy and the processing has been assessed in accordance with that policy.

c. Authentication credentials, API keys, secrets, and security configuration must not be submitted to any AI system under any circumstances.

5.3 Provider Due Diligence

a. AI providers are vendors and must be tiered and assessed in accordance with the Vendor Management Policy. Any AI provider processing Restricted or Confidential information is a Tier 1 vendor.

b. In addition to standard vendor due diligence, the following must be established and documented before an AI provider is approved:

i. Whether inputs submitted to the service are retained by the provider, and for how long.

ii. Whether inputs are used to train or improve the provider's models. Services that use CredX or customer inputs for provider model training must not be approved for Restricted or Confidential information.

iii. The geographic locations in which processing and storage occur.

iv. The provider's sub-processors.

v. Whether a Data Processing Agreement is required and has been executed.

c. Provider terms must be re-verified at least annually as part of the Tier 1 vendor review.

5.4 Customer Data and Model Training

a. Customer data must not be used to train or fine-tune any model, whether operated by CredX or by a provider, unless expressly permitted by the applicable customer agreement and approved by the Founder & CEO.

b. Where AI is used to enrich or process customer records, processing must be limited to inference. Customer data submitted for inference must not persist in any model.

c. Where a model or derived dataset is developed using data originating from more than one customer, it must not expose, or permit inference of, one customer's confidential information to another. Where this cannot be assured, processing must be isolated per customer.

d. AI processing of customer data must be disclosed to affected customers, and the AI provider recorded in CredX's sub-processor disclosures.

5.5 Automated Credit Decisioning Models

a. CredX's decisioning engine aggregates credit, transaction, loyalty, behavioural, and open banking data and applies machine learning models to assess risk and generate pre-approved offers or credit eligibility. These models constitute **automated decision-making about an individual's access to credit**, and are subject to the requirements in this section in addition to all other requirements of this policy.

5.5.1 Model Documentation and Inventory

a. Every model used in or influencing a credit decision must be recorded in the AI inventory with: its purpose, the decision it informs, its owner, its current version, the date it entered production, and the date of its last review.

b. Model documentation must record the features used, the provenance and time period of the training data, the target variable, known limitations, and the performance metrics against which the model was validated.

c. Models must be version-controlled. The model version that produced a given decision must be recoverable for any decision, for the duration of the applicable retention period.

5.5.2 Fairness and Disparate Impact

a. Models used in credit decisioning must be tested for disparate impact across protected characteristics before deployment, and at least annually thereafter.

b. Protected characteristics and close proxies for them must not be used as model features. Where a feature is found to act as a proxy, its use must be reassessed and documented.

c. Testing results, including any disparate impact identified and the remediation applied, must be retained as evidence and made available to the lending partner and to auditors on request.

d. A model that cannot be shown to meet the fairness requirements of this section must not be deployed into a credit decision.

5.5.3 Explainability and Adverse Decisions

a. The reasons for a credit decision must be explainable to the applicant in plain language, and in sufficient detail for the lending partner carrying regulatory accountability for that decision to substantiate it.

b. A model whose outputs cannot be explained to the standard above must not be the determining factor in an adverse decision.

c. Applicants subject to an adverse automated decision must be able to request human review. The availability of human review must be communicated at the point the decision is delivered.

d. Records of adverse decisions, the model version and inputs that produced them, and the outcome of any human review must be retained in accordance with the Data Retention Policy.

5.5.4 Consent Boundaries in Personalization

a. Behavioural, loyalty, transaction, and open banking data may be used only for the purposes the individual consented to. Data collected for decisioning must not be silently repurposed for personalization, targeting, or merchant analytics, and vice versa.

b. Where consent is withdrawn, the affected data must cease to be used as a model input, and downstream copies must be addressed under the retention and deletion requirements of the Privacy Management Policy.

c. Merchant-facing analytics derived from consumer data must be de-identified before disclosure, and must not be presented at a granularity that permits re-identification.

5.5.5 Monitoring and Change Control

a. Model performance must be monitored in production for drift, degradation, and anomalous decision patterns.

b. A change to a model, its version, its features, or its training data is a change to production behaviour affecting customers, and must be managed under the System Change Policy with the approvals required by this section.

c. Any use of a third-party model or scoring service within the decisioning path makes that provider a Tier 1 vendor under the Vendor Management Policy, and the provider's own model documentation and fairness evidence must be obtained.

5.5.6 Regulatory Accountability

a. The lending partner or sponsor bank retains regulatory accountability for lending decisions made on its charter, regardless of whether CredX's model produced the recommendation. CredX must be able to supply, on request and within the timeframes set in the partner agreement, the model documentation, decision records, and audit trails that partner requires to discharge that accountability.

b. The status of these controls must be represented accurately in due diligence. Where model documentation, fairness testing, or explainability mechanisms are specified but not yet implemented, that must be stated plainly rather than implied to be in place.

5.6 Human Oversight and Output Validation

a. AI systems must not make or materially influence decisions affecting customers, individuals, pricing, access rights, or transactions without human review. For automated credit decisioning, this requirement is met through the adverse-decision human review mechanism required by the Automated Credit Decisioning Models section, not by review of every decision.

b. AI-assisted output delivered to customers or written to production records is subject to the output controls of the Processing Integrity Policy. Output must be validated against specifications, and anomalous or unexpected output flagged for review.

c. Where AI is used to generate or enrich record attributes, prompts must constrain the system to information present in the source material, and generated attributes must be identifiable as AI-derived so that their provenance can be established.

d. Personnel remain accountable for any AI-assisted output they rely upon, publish, or commit.

5.7 AI-Assisted Software Development

a. AI coding assistants may be used only where approved under this policy.

b. Source code, configuration, and documentation authored with AI assistance are subject to the same peer review, testing, and deployment controls as all other changes under the Software Development Policy and System Change Policy. AI assistance does not reduce or replace any review requirement.

c. Credentials, secrets, production data, and customer data must not be submitted to AI coding assistants.

d. Dependencies introduced through AI-assisted development are subject to the same dependency and supply chain scanning requirements as all other dependencies.

5.8 Secure Integration

a. AI integrations must authenticate using credentials held in approved secrets management and must not embed credentials in source code.

b. All communication with AI providers must be encrypted in transit in accordance with the Encryption Policy.

c. AI integrations must fail closed. Where a provider is unavailable or returns an error, the system must return no result rather than proceed with unvalidated output.

d. Requests to AI providers must apply timeouts, and failures must be logged for monitoring and review.

e. Data submitted to AI providers must be minimized to that required for the defined purpose, with non-public fields excluded before transmission where the purpose does not require them.

5.9 Change Management

a. Changes to the model, model version, provider, prompt content, or the data submitted to an AI system are changes to production behaviour and must be managed in accordance with the System Change Policy.

5.10 Risk Assessment

a. AI-related risks must be assessed as part of the annual risk assessment conducted under the Risk Assessment and Management Policy, including risks of inaccurate output, data leakage to providers, provider dependency, and regulatory exposure.

b. A privacy assessment must be completed under the Privacy Management Policy before introducing any AI processing of personal information.

6 Roles and Responsibilities

a. The **Chief Technology Officer (CTO)** approves AI systems for use, maintains the AI inventory and model documentation, performs provider due diligence, commissions fairness and disparate impact testing, monitors deployed models for drift, and ensures AI integrations meet the requirements of this policy.

b. The **Founder & CEO** approves any use of customer data for model training, approves deployment of any model into a credit decision, and approves any exception to this policy.

c. **Operations** tracks the regulatory and consent obligations attaching to automated decisioning, and coordinates the supply of model documentation and decision records to lending partners and regulators.

d. **Managers** are responsible for ensuring their teams use only approved AI systems and understand the data restrictions in this policy.

e. **All employees and contractors** are responsible for complying with this policy, for validating AI-assisted output they rely upon, and for reporting suspected misuse or data exposure to the CTO.

7 Training

a. Acceptable use of AI, including the data restrictions defined in this policy, must be included in annual security awareness training under the Security Awareness Training Policy.

8 Enforcement

a. Violations of this policy may result in disciplinary action up to and including termination of employment or contract.

b. Submission of Confidential or personal information to an unapproved AI system must be reported and handled as a security incident in accordance with the Security Incident Response Policy.