Artificial Intelligence Governance Policy

IronHub Inc

August 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
aug 20 2026	Initial document

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 IronHub, 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 IronHub, including generative AI services, machine learning models, AI-assisted development tools, and AI features embedded within third-party software.

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

2 Background

a. AI systems introduce risks that are not fully addressed by conventional application controls. Information submitted to an AI service leaves IronHub'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. IronHub'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 IronHub 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 IronHub, 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 IronHub 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 IronHub or customer information. Personal or trial accounts with AI services must not be used for IronHub 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 IronHub 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 IronHub or by a provider, unless expressly permitted by the applicable customer agreement and approved by the President & 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 IronHub's sub-processor disclosures.

5.5 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.

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 equipment 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.

5.6 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.

5.7 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.

5.8 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.

5.9 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, performs provider due diligence, and ensures AI integrations meet the requirements of this policy.

b. The **President & CEO** approves any use of customer data for model training and any exception to this policy.

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

d. **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.