



Artificial Intelligence (AI)
Acceptable Use and
Governance Policy

Table of Contents

- 1. Purpose..... 3
- 2. Scope..... 3
- 3. Guiding Principles 3
- 4. Permitted Uses of AI..... 3
- 5. Prohibited Uses of AI..... 4
- 6. Data Protection and Confidentiality..... 4
- 7. Accuracy, Verification, and Accountability 5
- 8. Records Management and Transparency 5
- 9. Approval and Oversight..... 5
- 10. Training and Updates 6
- 11. Effective Date 6
- Appendix A: Approved AI Platforms – Rationale..... 6
- Appendix B: Staff Acknowledgment of AI Acceptable Use Policy..... 7
- Appendix C: Definitions 7

1. Purpose

The purpose of this policy is to establish clear standards for the responsible, ethical, secure, and compliant use of Artificial Intelligence (AI) tools by staff of the Staunton Redevelopment and Housing Authority (SRHA) and its affiliated entity, the Staunton Housing Corporation (SHC). This policy is intended to safeguard confidential information, ensure regulatory compliance, preserve public trust, and promote the appropriate use of AI as a productivity and decision-support tool—not as a substitute for professional judgment.

2. Scope

This policy applies to:

- All SRHA and SHC employees, officers, contractors, interns, and consultants;
- All uses of generative AI, machine learning, large language models, and similar automated tools (e.g., text generation, data analysis, summarization, image generation);
- All AI tools accessed through web-based platforms, software applications, or embedded systems, whether free or paid.

3. Guiding Principles

AI may be used by SRHA staff only in a manner that:

- Supports, but does not replace, human decision-making;
- Protects resident privacy and confidential information;
- Complies with HUD, state, and federal laws and regulations;
- Maintains transparency, accountability, and defensibility in agency actions;
- Preserves the integrity of official records and determinations.

4. Permitted Uses of AI

Subject to the restrictions in this policy, staff may use approved AI tools for the following purposes:

- Drafting and editing non-final documents (e.g., internal memos, policies, emails, grant narratives, meeting agendas);
- Summarizing publicly available regulations, guidance, or reports;
- Brainstorming ideas, outlines, or alternative language;
- Improving clarity, grammar, or organization of staff-prepared content;
- Generating templates or checklists for internal use.

All AI-generated content must be reviewed, validated, and approved by the staff member prior to use or dissemination.

5. Prohibited Uses of AI

The following uses of AI are strictly prohibited:

- Inputting or uploading any personally identifiable information (PII), protected health information (PHI), or sensitive resident data (including names, SSNs, DOBs, addresses, income details, or case records);
- Inputting or uploading confidential, non-public, or restricted information, including:
 - Resident files or eligibility determinations;
 - HUD systems data (e.g., IMS/PIC, EIV);
 - EIV (Enterprise Income Verification) data may never be entered into any AI tool under any tier or circumstance; EIV is governed by separate HUD security requirements that carry independent penalties;
 - Procurement-sensitive information;
 - Privileged legal communications or attorney work product, except as permitted under Section 6;
 - Personnel records;
- Using AI to make or automate eligibility determinations, rent calculations, compliance decisions, or enforcement actions;
- Representing AI-generated output as official SRHA or SHC determinations without independent verification;
- Using AI tools to circumvent established policies, controls, or approval processes.

6. Data Protection and Confidentiality

- Staff shall not assume that AI platforms are secure or confidential.
- Information entered into AI tools may be retained, logged, or used to train models outside the control of SRHA.
- On consumer, free, or otherwise unapproved tiers, only de-identified, hypothetical, or publicly available information may be used in AI prompts. On AI tools approved under Section 9, non-public SRHA or SHC business information may be used; however, personally identifiable information (PII), protected health information (PHI), and resident data remain prohibited unless a Business Associate Agreement or equivalent data-protection agreement is in place and use has been expressly authorized in writing by the Executive Director.
- Any uncertainty regarding whether information is permissible to use with AI shall be resolved in favor of non-use.
- Any staff member who enters prohibited information into an AI tool, or who suspects a data exposure involving an AI tool, shall immediately report the incident to the Executive Director or designee and follow the agency's established data-breach and incident-response procedures, including applicable federal and state breach-notification requirements.
- Privileged communications and attorney work product shall not be entered into AI tools on consumer or unapproved tiers. On AI tools approved under Section 9, such material may be used to analyze, summarize, or pressure-test legal questions only with the express

authorization of the Executive Director and, where practicable, the awareness of counsel. Because the application of the attorney-client privilege and work-product protection to AI tools is unsettled, staff shall weigh the risk that submitting such material to any third-party tool could affect those protections, and shall consult counsel where the risk is material.

7. Accuracy, Verification, and Accountability

- AI-generated content may contain errors, omissions, outdated information, or fabricated citations.
- Staff are solely responsible for verifying the accuracy, completeness, and appropriateness of any AI-assisted work product.
- AI output shall never be relied upon as the sole basis for:
 - Regulatory interpretation;
 - Financial decisions;
 - Compliance determinations;
 - Resident eligibility or rent calculations.

8. Records Management and Transparency

- AI-generated drafts that become part of the official record must comply with SRHA records retention requirements.

9. Approval and Oversight

- The Executive Director retains authority to approve, restrict, or prohibit the use of specific AI tools.
- At this time, the **only approved AI service providers for SRHA and SHC staff use are OpenAI and Anthropic**, and only through their commercial, business, enterprise, or API tiers as defined and conditioned below.
- An AI service may be approved for SRHA or SHC business use only if it satisfies all of the following conditions:
- It is a paid commercial, business, enterprise, or API tier of an approved provider, not a free or consumer or individual tier, operating under written terms accepted on behalf of SRHA;
- It contractually excludes SRHA and SHC inputs and outputs from use in training or improving AI models, by binding contract rather than by a user-adjustable setting; and
- It operates under a Zero Data Retention agreement where the provider offers one, together with contractual data-minimization terms, so that personally identifiable information is not retained, logged for profiling, or otherwise tracked beyond the minimal abuse-monitoring the provider requires; where Zero Data Retention is unavailable, a defined maximum retention period must be specified in writing.
- **Only AI tools that are vetted, approved, and made available by SRHA may be used for SRHA or SHC business purposes.** Staff may not independently select, download, subscribe to, or configure AI tools for agency use.

- Before any AI tool is approved, it shall be vetted against the criteria in this Section, including executed data-protection terms (such as a Data Processing Addendum or Business Associate Agreement where applicable), and its procurement shall comply with the agency’s procurement policy and applicable federal requirements, including 2 CFR Part 200 where federal funds are used.
- The Executive Director serves as the AI administrator, responsible for maintaining the approved-tools list, coordinating vetting, and tracking acknowledgments, and may delegate these functions in writing to the Deputy Director.
- **Browser extensions, plug-ins, third-party integrations, embedded AI features in other software, or automated agents are prohibited** unless expressly approved in writing by the Executive Director.
- This prohibition is directed at AI features, agents, and services that transmit SRHA or SHC data to systems not approved under this Section. AI features embedded within enterprise software already approved by SRHA, and governed by that software’s existing data-protection agreement, are permitted within the scope of that agreement.
- Use of any other AI tools, platforms, browser extensions, or embedded AI services requires prior written approval from the Executive Director.
- SRHA may maintain and periodically update a list of approved or disapproved AI platforms.
- Misuse of AI tools or use of unapproved platforms may result in corrective action, up to and including disciplinary measures.

10. Training and Updates

- Staff may receive training or guidance on appropriate AI use as technologies evolve.
- This policy will be reviewed periodically and updated as necessary to reflect regulatory changes, technological developments, or agency needs.

11. Effective Date

This policy is effective upon issuance and applies to all AI use going forward.

Appendix A: Approved AI Platforms – Rationale

SRHA has limited approved AI use to OpenAI and Anthropic based on the following considerations:

- Market leadership and stability in generative AI technologies;
- Publicly documented security, privacy, and governance practices;
- Availability of enterprise-grade controls and transparency regarding model behavior;
- Widespread adoption across regulated industries;

- Availability of executed data-protection terms, including contractual no-training commitments, defined data retention, and Zero Data Retention or Business Associate Agreement options where required.

Approval of these providers does not imply endorsement of any specific output and does not reduce staff responsibility for verification, judgment, or compliance.

Appendix B: Staff Acknowledgment of AI Acceptable Use Policy

All SRHA and SHC employees, officers, contractors, interns, and consultants authorized to use AI tools must, before being granted access and upon each material revision of this policy, acknowledge the following:

I understand that OpenAI and Anthropic are the only approved AI platforms for use in connection with SRHA and SHC work. I agree to comply with the Artificial Intelligence Acceptable Use and Governance Policy, including restrictions on confidential information, resident data, and decision-making. I understand that misuse of AI tools or use of unapproved platforms may result in corrective or disciplinary action.

Employee Name: _____

Signature: _____

Date: _____

Adopted by the Executive Director, Staunton Redevelopment and Housing Authority, and ratified by the Board of Commissioners on _____

Appendix C: Definitions

Artificial Intelligence (AI) / Generative AI. Software that uses machine learning or large language models to generate text, images, code, analysis, or other outputs in response to prompts.

Automated Agent. An AI tool configured to take actions, access systems, or complete multi-step tasks with limited or no human review of each step.

Personally Identifiable Information (PII). Information that identifies or can be used to identify an individual, including names, Social Security numbers, dates of birth, addresses, income details, and case records.

Protected Health Information (PHI). Individually identifiable health information, including reasonable-accommodation, disability, and medical records held by the agency.

Zero Data Retention (ZDR). A contractual arrangement under which a provider does not store inputs or outputs beyond what is minimally required for abuse monitoring.

Consumer / Individual Tier and Commercial Tier. A consumer or individual tier is a free or personal paid subscription marketed to individuals (e.g., free or pro consumer plans). A commercial tier is a business, enterprise, or API offering provided under negotiated terms that contractually exclude customer data from model training; only commercial tiers may be approved under Section 9.