

BETA SYSTEMS INTERNATIONAL

Sample Artificial Intelligence Use and Governance Policy

Illustrative policy framework for a California public housing authority

Proposal context	Prepared for consideration	Version
2026-RFP-DP-003 Consulting Services to Advise on Artificial Intelligence Use	Housing Authority of the City of Alameda	Sample 1.0 September 17, 2026

Purpose of this sample

This proposal exhibit demonstrates a practical policy architecture for moving from a near-prohibition on workplace AI to controlled, mission-aligned use. It is tailored to a small public agency that administers housing programs, develops and manages affordable housing, delivers social services, and holds substantial personally identifiable and confidential information.

Important status statement

This is a newly prepared illustrative sample. It is not represented as a policy previously adopted by AHA or another government agency. Final policy language should be validated against AHA’s systems, collective bargaining obligations, records schedules, insurance, contracts, legal advice, and operational procedures before adoption.

Prepared by BETA Systems International

Proposal support document | Government AI governance and implementation

Executive Policy Summary

AHA can enable useful AI without exposing client data or delegating public authority to a machine. This model establishes an approval path, separates low-risk assistance from high-impact decisions, requires human review, limits data inputs, governs vendors, and creates an auditable inventory of every approved use.

Policy outcome	Control
Permit bounded productivity use	Only approved tools and approved use cases; public or low-sensitivity information by default
Protect applicants, participants, tenants, landlords, employees and partners	No confidential or regulated data in public AI tools; least privilege, minimization, retention and incident controls
Preserve lawful human authority	AI may assist but may not make final consequential decisions or communicate unverified determinations
Reduce discrimination and accessibility risk	Impact review, representative testing, appeal paths, language access and accessible alternatives
Create management visibility	AI inventory, risk tier, owner, approval, monitoring metrics, change control and annual review
Control supplier risk	Mandatory disclosure, security and privacy review, audit rights, data-use limits, exit and deletion terms

Recommended Initial Posture

- Start with a 90-day controlled pilot using low-risk internal tasks such as drafting, summarization of public material, meeting agendas, job aids, translation drafts, and code assistance with review.
- Keep Housing Choice Voucher eligibility, rent calculation, termination, reasonable accommodation, tenant selection, fraud referral, employee action, procurement award, and other rights-affecting uses prohibited unless separately approved as high-risk systems with legal review and independent validation.
- Do not enter Social Security numbers, dates of birth, disability or medical details, immigration information, financial account data, criminal history, case notes, VAWA information, credentials, or unredacted client records into a general-purpose AI service.
- Designate an AI Governance Lead and a small cross-functional review group rather than creating a large new bureaucracy.

Policy Administration

Field	Model entry
Policy owner	Executive Director or designee
Operational steward	AI Governance Lead

Field	Model entry
Required reviewers	Information security, privacy, legal, records, accessibility/civil rights, program owner, HR or procurement as applicable
Applies to	Employees, temporary staff, interns, commissioners when acting for AHA, contractors, consultants, vendors and service providers
Review cycle	At least annually and after material legal, technology, vendor, incident or mission changes
Supersedes	Any inconsistent informal AI practice; stricter law, contract, grant, security or program rule controls

1 Purpose and Policy Statement

AHA may use artificial intelligence when the use advances its mission, is lawful and necessary, produces a documented benefit, and is proportionate to the risk. AI is an assistive capability. It does not replace accountable public officials, program expertise, due process, professional judgment, or the obligation to provide accurate and equitable service.

AHA will govern AI throughout its lifecycle: proposal, assessment, procurement or configuration, testing, approval, operation, monitoring, change, incident response and retirement. No employee or supplier may deploy an AI capability on AHA's behalf outside this process.

2 Scope

This policy applies to stand-alone AI tools and AI features embedded in software, websites, search, productivity suites, customer service tools, security products, analytics, case management, property management, finance, human resources, development, communications and vendor services. It covers free, trial, personally purchased and agency-provided accounts when used for AHA work.

Conventional deterministic automation may be managed under existing IT controls, but it enters this policy when it learns from data, infers, predicts, recommends, generates content, ranks people or cases, uses biometric or natural-language functions, or materially influences a decision.

3 Definitions

Term	Policy meaning
Artificial intelligence or AI	A machine-based system that generates predictions, recommendations, decisions or synthetic content for human-defined objectives.
Generative AI	AI that creates or transforms text, images, audio, video, code or other content.

Term	Policy meaning
AI system	The model plus data, prompts, interfaces, integrations, people, procedures and downstream processes that make up the operational use.
Consequential decision	A decision or recommendation that affects eligibility, benefits, housing, tenancy, accommodations, employment, procurement, legal rights, safety or access to essential services.
High-impact use	An AI use that makes, recommends, materially supports or communicates a consequential decision, or processes sensitive data at scale.
Human review	Meaningful review by a trained person with time, authority, context and ability to disagree, correct, stop and document the outcome.
Sensitive information	PII, confidential program data, protected or security information, credentials, financial information, health/disability information, VAWA information and any data restricted by law, contract or policy.
Public AI tool	A consumer or generally available service without an AHA-approved enterprise agreement and security/privacy configuration.
AI inventory	The authoritative register of proposed, piloted, active, suspended and retired AI use cases and systems.

4 Governing Principles

1. Mission and necessity. Begin with a defined public-service problem; do not use AI merely because it is available.
2. Human accountability. AHA personnel remain responsible for decisions, communications, records and outcomes.
3. Privacy and minimization. Use the least data needed, prefer de-identified or synthetic data, and do not reuse data beyond the approved purpose.
4. Security and resilience. Apply defense in depth, least privilege, testing, logging, change control and safe failure modes.
5. Fairness and civil rights. Evaluate disparate impacts, data quality, accessibility, language access and reasonable accommodation.
6. Transparency and contestability. Provide proportionate notice and a practical path to ask questions, correct data, request human review and appeal.
7. Accuracy and evidence. Treat outputs as unverified until checked against authoritative sources.
8. Procurement accountability. Require vendors to disclose AI and accept enforceable data, security, audit and change obligations.
9. Lifecycle governance. Monitor performance and retire uses that no longer meet law, need, accuracy, equity, cost or security expectations.

5 Roles and Responsibilities

Role	Core responsibilities
Board or governing authority	Receives material risk reporting and approves policy or high-impact uses when required by governance rules.
Executive Director	Sets risk tolerance, designates accountable officials, approves high-risk exceptions and ensures resources.
AI Governance Lead	Maintains inventory, coordinates reviews, records approvals, monitors changes, schedules training and reports metrics.
AI Review Group	Cross-functional review by program, IT/security, privacy, legal, records, accessibility/civil rights, procurement and HR as needed.
System or use-case owner	Documents need, data, workflow, testing, human review, performance, complaints, changes and retirement.
Information security and privacy	Classifies data, evaluates architecture and vendor controls, approves integrations, monitors events and supports incident response.
Legal and records officials	Review authority, due process, civil rights, labor, public records, retention, discovery, intellectual property and contract terms.
Managers	Authorize staff access, reinforce approved uses, review work products and report misuse or drift.
Users	Use only approved tools and purposes, protect data, verify outputs, disclose use when required and report errors or incidents.
Vendors and contractors	Meet the same or stricter controls, disclose model and subprocessors, support audits, notify changes/incidents and return or delete data.

6 Risk Classification

The AI Governance Lead assigns a tier based on the highest applicable factor. A use does not become low risk merely because a vendor describes it as an assistant or because a human clicks the final button.

Tier	Typical examples	Approval and control
Tier 1 Low	Drafting from public information; brainstorming; formatting; summarizing public reports; generating nonproduction code	Manager-approved catalog use; trained user; no sensitive data; verification required
Tier 2 Moderate	Internal knowledge search; translation draft; meeting transcription; analytics using internal operational data; constituent chatbot limited to published information	Documented assessment; program owner plus IT/privacy approval; testing, logging, notice where public-facing

Tier	Typical examples	Approval and control
Tier 3 High	Eligibility or rent support; termination; reasonable accommodation; tenant screening; fraud prioritization; employee screening; procurement evaluation; safety decisions	Presumptively restricted; executive, legal, security, privacy and civil-rights review; impact assessment; independent validation; appeal and continuous monitoring
Prohibited	Autonomous final decisions; covert surveillance; social scoring; deceptive impersonation; unlawful discrimination; bypass of due process; unapproved biometric identification; use of confidential data in public tools	May not be used. Exceptions cannot waive law, civil rights, due process or security obligations.

7 Acceptable Use Rules

7.1 Before use

- Confirm that the tool, account, feature and proposed task appear in the approved AI catalog or obtain approval.
- Use the minimum necessary information. Remove names, identifiers and case details whenever the task can be completed without them.
- Understand whether prompts, files and outputs are retained, used for model training, accessible to administrators, or transferred to subprocessors.
- For public communications or operational work, identify the authoritative source against which the output will be checked.

7.2 During use

- Do not treat an AI response as fact, policy, legal advice, a case determination or a complete record.
- Do not connect tools to email, shared drives, case systems, calendars, code repositories or databases without integration approval and scoped access.
- Do not disable logging, safety, retention, identity or security controls.
- Do not ask AI to imitate a real person or create misleading media, signatures, testimony, evidence or records.

7.3 Before reliance or release

- Verify facts, calculations, citations, quotations, policy language, translations and code. Open cited sources rather than trusting a generated citation.
- Check for protected-class bias, inaccessible content, culturally inappropriate language and missing context.
- Edit the output so it reflects AHA's judgment and voice. The responsible employee remains the author or reviewer of record.
- Preserve prompts, source materials, validation and approvals when required by the records schedule, litigation hold, audit or use-case plan.

8 Data Protection and Privacy

AHA data classifications and existing privacy and security requirements apply regardless of whether information is typed, uploaded, recorded, inferred, embedded, retrieved through a connector, or included in an output.

Data category	Default AI rule
Public information	Permitted in approved tools for approved purposes, subject to accuracy and records controls.
Internal operational information	Only in approved enterprise tools when the use, access, retention and vendor terms are documented.
Confidential or regulated information	Do not use unless the specific system and use case have written approval, a lawful basis, minimum-necessary fields, access controls and contract protections.
Highly sensitive information	Social Security numbers, credentials, payment data, detailed health/disability information, VAWA information, criminal history and similarly sensitive records are prohibited in general-purpose AI and require exceptional written authorization in a purpose-built environment.
De-identified or synthetic data	Preferred for testing, but re-identification risk and realism must be evaluated; synthetic data must not be represented as an actual person or event.

8.1 Required safeguards

- Data minimization and field-level restriction; no bulk copying when a smaller excerpt will work.
- Enterprise identity, multifactor authentication, role-based access and periodic access review.
- Encryption in transit and at rest; tenant isolation; secure key and secret management.
- Contractual prohibition on vendor training, sale, advertising use or unrelated secondary use of AHA data unless expressly authorized.
- Defined retention and deletion for prompts, files, outputs, logs, embeddings, caches and backups.
- Security logging sufficient to investigate access, changes, integrations and incidents without creating unnecessary sensitive logs.
- Privacy review when combining datasets, creating profiles, inferring sensitive traits or monitoring individuals.

9 Human Oversight and Consequential Decisions

AI shall not be the final decision-maker for any action affecting a person's housing, assistance, tenancy, accommodation, employment, contracting opportunity, legal rights or access to services. Human review must be substantive, not ceremonial.

Required element	Minimum standard
Qualified reviewer	Understands the program, governing rules, system limits and signs of unreliable output.

Required element	Minimum standard
Authority	Can reject the output, obtain more evidence, correct data, pause the process and escalate.
Independent evidence	Reviews the underlying record and governing criteria, not only the AI score or summary.
Explanation	Can state the actual facts and rules supporting the decision in understandable terms.
Notice and recourse	Affected person receives applicable notice, contact information, correction process, reasonable accommodation and appeal or hearing rights.
Documentation	Records AI's role, reviewer, evidence checked, changes made and final rationale when the use-case plan requires it.

10 Civil Rights Equity Accessibility and Language Access

AI use must comply with applicable fair housing, civil rights, disability, language access, employment and nondiscrimination duties. Technology cannot narrow a protection, shift the burden to an applicant or participant, or make an inaccessible channel the only practical route to service.

- Test outcomes and error rates across relevant groups and circumstances using lawful, representative methods. Aggregate testing must not create a new basis for individual profiling.
- Evaluate proxy variables, historical bias, missing data, feedback loops and whether false positives or false negatives harm some groups more severely.
- Provide a non-AI channel and qualified human assistance for people who cannot or do not wish to use an AI interface when rights or essential services are involved.
- Validate translations and plain-language explanations, especially for vital documents, deadlines, eligibility requirements, adverse actions and appeal rights.
- Test public-facing tools for keyboard access, screen-reader compatibility, contrast, focus order, captions, timeouts and error recovery under AHA's accessibility standard.

11 Transparency Records and Public Communications

11.1 Notice

AHA will provide notice proportionate to impact. A public-facing chatbot must identify itself as automated, describe its scope, warn users not to submit restricted information, link to privacy information, and offer a human channel. For a consequential process, notice must also explain AI's role, principal data sources, human review, correction and appeal routes, subject to lawful security and confidentiality limits.

11.2 Records management

AI content created or received in the conduct of AHA business may be an agency record regardless of device, account or medium. The program owner and records official will determine retention based on content and function. Users must not rely on a vendor's chat history as the official recordkeeping system.

- Store official work products and required supporting records in approved AHA repositories.
- Apply legal holds, audit requirements and public-records review to responsive AI records.
- Separate transitory experimentation from records supporting decisions, policies, notices, contracts or program administration.
- Do not disclose confidential, privileged, security-sensitive or legally exempt material merely because it appears in an AI output.

12 Procurement and Vendor Management

All solicitations and vendor intake processes should require disclosure of intentional and embedded AI. The disclosure obligation continues throughout the contract because vendors may add AI after award.

Due diligence area	Required evidence or term
System description	Models, purpose, decision role, data flows, integrations, hosting locations, subprocessors and whether models can change automatically.
Data rights	AHA ownership/control; purpose limitation; no model training or sale; return/export; verified deletion; rules for derived data and embeddings.
Security	Independent assurance, vulnerability management, encryption, access controls, logs, secure development, incident response and notification.
Performance	Defined metrics, limitations, test results, subgroup analysis where relevant, drift monitoring and correction process.
Transparency	Documentation sufficient for AHA to explain use, manage records, answer complaints and respond to audits or lawful requests.
Change control	Advance notice and approval for material model, provider, training-data, feature, integration, retention, location or subprocessor changes.

Due diligence area	Required evidence or term
Accountability	Audit and assessment rights, cooperation, indemnity/insurance as appropriate, service levels, suspension rights and termination assistance.
Accessibility and language	Conformance evidence, remediation commitments, supported languages and human escalation.
Exit readiness	Data portability, configuration export, continuity, deletion certificate and avoidance of proprietary lock-in.

13 AI Intake Assessment and Approval

No pilot or production use begins until the minimum review for its tier is complete. Review is based on the operational system, not marketing claims about a model.

Gate	Decision evidence
1 Define	Problem, intended users, expected public or operational benefit, alternatives and success measures.
2 Classify	Decision impact, affected people, data categories, scale, reversibility and risk tier.
3 Assess	Privacy, security, civil rights, accessibility, records, legal, labor, procurement, fiscal and operational impact.
4 Test	Representative scenarios, unacceptable outputs, accuracy, hallucination, injection, access boundaries, bias, accessibility and fallback.
5 Approve	Named owner, authorized users, conditions, training, notice, monitoring, expiration and approver signatures.
6 Operate	Metrics, sampling, complaints, incidents, vendor notices, costs, overrides and periodic review.
7 Change or retire	Reassess material changes; suspend when thresholds fail; export records and verify deletion at retirement.

14 Testing Monitoring and Change Control

Testing shall reflect the actual workflow, data conditions and affected population. A successful demonstration is not production validation. High-impact uses require an evaluation plan independent enough to challenge the system owner and vendor.

Control area	Illustrative metric
Quality	Factual accuracy, citation validity, completeness, calculation accuracy and reviewer correction rate.
Equity	Error and outcome differences across relevant populations; accommodation and language access success.
Safety and privacy	Restricted-data events, prompt-injection success, unauthorized access, leakage and harmful-content rate.
Human oversight	Override rate, review time, escalations, appeals, reversals and evidence of automation bias.
Service	Resolution, wait time, abandonment, human handoff, accessibility defects and user complaints.
Operations	Availability, latency, cost per task, vendor changes, model drift and incident recovery.

The system owner shall define thresholds before launch. AHA may pause a use immediately when law, safety, security, fairness, accuracy, accessibility, cost or public trust is at risk. Material changes require reassessment before release.

15 Security and Incident Response

AI incidents follow AHA's security, privacy, records, legal and operational response processes, with AI-specific facts added to triage. Users must report suspected exposure, harmful output, unauthorized integration, impersonation, manipulation, discriminatory outcome, vendor change, or uncontrolled automated action immediately through the designated channel.

Phase	AI-specific action
Contain	Disable access or integration; preserve evidence; stop automated actions; protect affected people and services.
Assess	Identify prompts, files, outputs, accounts, logs, model/version, connectors, downstream recipients, data classes and affected decisions.
Notify	Engage security, privacy, legal, records, program and leadership; meet contractual and legal notices; communicate accurately.
Correct	Revoke access, delete exposed data where possible, correct records and decisions, contact affected persons when required, retrain users and fix controls.

Phase	AI-specific action
Learn	Document cause, vendor response, impacts and corrective actions; update risk assessment, tests, contracts, policy and training.

16 Training and Workforce Practices

- All personnel receive foundational training before access and annually thereafter: approved tools, data restrictions, verification, records, bias, accessibility, security and incident reporting.
- Managers and use-case owners receive added training on risk classification, human oversight, testing, monitoring and vendor management.
- Technical and procurement personnel receive role-specific training on architecture, model/vendor changes, threat testing, contract protections and evaluation evidence.
- AHA will consult applicable labor agreements and workforce processes. AI will not be introduced covertly to monitor workers or make employment decisions.
- Training completion and policy acknowledgment are recorded. Access may be suspended when training expires.

17 Enforcement Exceptions and Reporting

Violations may result in removal of access, corrective action, contract remedies or discipline consistent with law, policy and labor obligations. Good-faith reporting of risk, error or misuse is encouraged and must not be retaliated against.

An exception request must identify the requirement, need, duration, affected data and people, alternatives, compensating controls, owner and approvers. Exceptions are time-limited and recorded. No exception may authorize unlawful discrimination, eliminate required due process, waive a person's rights, or permit an autonomous final consequential decision.

Appendix A Approved Use Case Record

The following fields form the minimum inventory and approval record. AHA may implement them in a service-management, governance, procurement or records system.

Field	Required entry
Identity	Use-case ID, name, department, owner, vendor, product, model/version and status.
Purpose	Problem, benefit, users, affected people, outputs, downstream action and non-AI alternative.
Data	Sources, fields, classification, legal authority, minimization, location, retention, deletion and sharing.
Risk	Tier, consequential-decision role, reversibility, scale, civil-rights/accessibility analysis and misuse scenarios.
Controls	Human review, access, notice, verification, logs, fallback, incident channel and prohibited actions.
Evidence	Security/privacy review, contract, test plan/results, vendor documents, accessibility evidence and training.
Decision	Approvers, conditions, pilot dates, production date, expiration, review cadence and suspension thresholds.
Operations	Metrics, complaints, appeals, overrides, incidents, costs, changes, reassessments and retirement evidence.

Appendix B Rapid User Checklist

- ☐ Approved? The tool, account, feature and task are authorized.
- ☐ Necessary? AI adds a defined benefit and a simpler non-AI method is not better.
- ☐ Safe data? The prompt and attachments contain no restricted information for this environment.
- ☐ Reliable? I can verify the result against an authoritative source.
- ☐ Fair and accessible? I checked for bias, language, accommodation and access barriers.
- ☐ Human accountable? A qualified person will make and explain any decision.
- ☐ Record handled? I know what must be saved and where.
- ☐ Transparent? I will disclose AI involvement when policy or context requires it.
- ☐ Reportable? I know how to stop and report a security, privacy, accuracy or harm concern.

Appendix C Illustrative Approval Matrix

Use case	Default tier	Default disposition
Draft a memo using public source material	1	Allow in approved tool with human verification
Summarize an internal meeting with no client case details	2	Assess recording, consent, access, retention and accuracy

Use case	Default tier	Default disposition
Translate a routine outreach flyer	2	Use qualified review for public release; vital documents require stronger validation
Answer website FAQs from an approved public knowledge base	2	Allow with scope limits, notice, logging, privacy warning and human handoff
Summarize an individual HCV case file	3	Restricted; purpose-built approved environment and substantive human review required
Recommend eligibility, rent, termination or fraud action	3	Presumptively restricted; never autonomous; full impact assessment and legal review
Rank job applicants or analyze recorded interviews	3	Presumptively restricted; employment, bias, notice and accommodation review
Generate deceptive voice/video of a real person	Prohibited	Do not use
Upload unredacted client records to a consumer chatbot	Prohibited	Do not use

Appendix D Implementation Roadmap

Period	Management actions	Deliverables
Days 0 to 30	Designate lead and review group; inventory current and embedded AI; suspend unknown high-risk uses; select approved pilot tools; establish intake and reporting.	Interim directive; inventory; risk rubric; approved-tool catalog; user notice; incident path
Days 31 to 60	Complete privacy/security/vendor reviews; train pilot users; test low-risk scenarios; define records handling and performance thresholds.	Assessments; pilot plan; training records; test report; contract controls
Days 61 to 90	Operate controlled pilot; sample outputs; track errors, costs, complaints and time saved; decide scale, revise or stop.	Pilot dashboard; lessons learned; policy revisions; production decision
Months 4 to 12	Expand by evidence; conduct high-impact discovery separately; integrate procurement disclosure and annual review; report to leadership.	Governance cadence; audited inventory; vendor change register; annual report

Appendix E Reference Framework

The final policy should be validated by counsel and AHA officials. This sample draws on the following authoritative frameworks and agency context available as of September 17, 2026:

1. National Institute of Standards and Technology, Artificial Intelligence Risk Management Framework 1.0, NIST AI 100-1, including the Govern, Map, Measure and Manage functions. <https://doi.org/10.6028/NIST.AI.100-1>
2. National Institute of Standards and Technology, Artificial Intelligence Risk Management Framework Generative Artificial Intelligence Profile, NIST AI 600-1. <https://doi.org/10.6028/NIST.AI.600-1>
3. State of California, GenAI Guidelines for Public Sector Procurement, Uses and Training, March 2024. <https://www.cdt.ca.gov/wp-content/uploads/2024/07/3a-GenAI-Guidelines.pdf>
4. California Department of Justice, Legal Advisory on the Application of Existing California Laws to Artificial Intelligence. <https://oag.ca.gov/system/files/attachments/press-docs/Legal%20Advisory%20-%20Application%20of%20Existing%20CA%20Laws%20to%20Artificial%20Intelligence.pdf>
5. U.S. Department of Housing and Urban Development, Fair Housing Act guidance concerning AI and tenant screening and advertising, May 2, 2024 archive. <https://archives.hud.gov/news/2024/pr24-098.cfm>
6. U.S. Department of Housing and Urban Development, Privacy Handbook, Revision 1.0. <https://www.hud.gov/sites/default/files/OCHCO/documents/privacy-handbook-2021.pdf>
7. U.S. Department of Housing and Urban Development, PII Minimization Plan. <https://www.hud.gov/sites/dfiles/OCHCO/documents/privacy-pii-minimization-plan.pdf>
8. Housing Authority of the City of Alameda, Privacy Policy Statement, effective September 27, 2022. https://www.alamedahsg.org/wp-content/uploads/2022/09/Privacy-Policy_final-draft.pdf
9. Housing Authority of the City of Alameda, public agency documents and program context. <https://www.alamedahsg.org/resource-library/agency-documents/>

Customization Notes for AHA

Before adoption, replace model entries with AHA's actual roles, data rules, records schedules, incident contacts, procurement and labor requirements, approval authorities and existing policies. Reconcile the final text with AHA's Administrative Plan, Language Access Plan, accommodation and VAWA procedures, privacy statement, enterprise systems, contracts and current AI policy.

Illustrative proposal exhibit • September 2026