GENERATIVE AI GOVERNANCE FRAMEWORK

Eight Pillars for Responsible AI Use in Local Government

Prepared by:

BoldPath CONSULTING

PREAMBLE

A MESSAGE TO THE GOVERNING BODY AND LEADERSHIP

Generative artificial intelligence refers to software that produces written text, summaries, analyses, and other content in response to instructions from a user. Tools of this kind are now widely available, including through many software platforms that local governments already use for day-to-day operations.

These tools offer real opportunities for local government. They can help employees draft documents, research questions, summarize information, and complete administrative tasks more efficiently. They also carry real risks. They can produce content that is confident but factually wrong. They can expose sensitive government information if not used carefully. And they can create confusion about who is responsible when AI-assisted work produces an error, a disclosure problem, or a decision that fails legal scrutiny.

This framework gives [Organization Name] a structured approach to capturing the benefits of generative AI while managing those risks. It is designed to be adopted by governing bodies that may not be technically expert, implemented by department heads who manage operational teams, and administered by central leadership with responsibility for human resources, information technology, records management, and legal compliance. The same document serves all three audiences, with each section clearly identifying who its primary guidance is directed to.

The framework is organized into eight pillars. Each pillar addresses a distinct governance question that every local government must answer in order to use AI tools responsibly. Together, the eight pillars constitute a complete AI governance policy that is appropriate for adoption as written, with bracketed placeholders filled in to reflect organizational specifics.

> **Important: Legal Review Required Before Adoption**
> Throughout this document, shaded boxes labeled "Legal Consideration" identify areas where legal requirements, obligations, or risks are relevant to implementation. These call-outs are designed to alert the organization and its leadership to issues that require legal analysis — they are not legal advice and should not be treated as such.
>
> Before adopting this policy, [Organization Name] should provide the complete document to its legal counsel and request review of all Legal Consideration call-outs. Legal requirements vary depending on the type of government entity, applicable state and federal law, collective bargaining agreements, and the organization's specific operations. Only legal counsel familiar with this organization's circumstances can advise on the specific obligations that apply.

This framework scales to organizational size and capacity. A small village or township implementing one AI tool requires less procedural infrastructure than a county operating multiple departments with diverse AI applications. The principles in each pillar apply to any organization; how deeply each pillar is implemented should reflect the organization's actual scale of AI use. Bracketed placeholders throughout the document indicate where organizational specifics are required.

GOVERNING BODY SUMMARY

EIGHT COMMITMENTS

By adopting this policy, [Organization Name] commits to the following principles governing the use of artificial intelligence in its operations:

1. **Pillar 1 - Purpose and Accountability:** We commit to ensuring that AI tools support our employees without replacing their judgment or their accountability for official work.
2. **Pillar 2 - Prohibited Uses:** We commit to identifying the decisions that may not be delegated to AI systems and requiring that qualified employees make those determinations directly.
3. **Pillar 3 - Approved Tools:** We commit to governing which AI tools may be used for official government work and maintaining oversight of those tools and the vendors that provide them.
4. **Pillar 4 - Data Protection:** We commit to protecting the confidentiality of the information our organization holds by establishing clear standards for what may and may not be entered into AI systems.
5. **Pillar 5 - Verification:** We commit to requiring that AI-generated content be reviewed and verified by qualified employees before it is relied upon in official work.
6. **Pillar 6 - Records and Disclosure:** We commit to ensuring that AI-assisted work products are identified, retained, and disclosed consistently with our public records and records management obligations.
7. **Pillar 7 - Workforce Readiness:** We commit to preparing our employees to use AI tools responsibly and to providing the training, oversight, and support necessary for accountable adoption.
8. **Pillar 8 - Governance and Oversight:** We commit to maintaining active oversight of AI use in this organization, reviewing our policies as technology and law evolve, and reporting AI governance activities to this governing body.

DEFINITIONS

KEY TERMS

The following terms are used throughout this framework. Definitions are written for a general audience and are intended to support consistent understanding and application of this policy.

Generative AI	Computer software that produces written text, summaries, analyses, images, or other content in response to instructions typed by a user. Common examples include tools that draft documents, answer research questions, or summarize information based on the user's request.
AI-generated content	Any text, summary, analysis, recommendation, or other output produced in whole or in part by a generative AI tool, regardless of how much the employee edited or revised it before use.
Prompt	The instruction or question a user types into an AI tool to produce a response. Prompts may include context, background information, or specific requests for format or content.

Output	The response produced by an AI tool in response to a prompt. Outputs must be reviewed and verified before being used in official work.
Hallucination	A term used to describe AI-generated content that is factually incorrect, fabricated, or misleading, even when it appears confident and plausible. AI tools can hallucinate citations, statistics, names, dates, and legal references.
Approved tool	A generative AI platform or application that has been reviewed and authorized by the organization for use in official work. Only approved tools may be used for government business.
Sensitive data	Information that is subject to legal protection, confidentiality obligations, or heightened privacy expectations. Examples include personnel records, medical information, criminal justice records, legal communications, and personally identifiable information.
Non-delegable decision	A government decision that requires direct human evaluation and judgment, and that may not be made solely on the basis of AI-generated output. Examples are identified in Pillar 2.
AI Governance Lead	The designated staff member is responsible for administering this policy, including maintaining the approved tool list, managing the escalation process, coordinating training, and tracking incidents. See Pillar 8.
Pillar	One of the eight governance domains addressed by this framework. Each pillar covers a distinct area of AI governance policy and includes guidance for the governing body, department heads, and central leadership.

PILLAR 1: Purpose and Human Accountability

COMPONENT 1: CORE PRINCIPLE

PRIMARY AUDIENCE: GOVERNING BODY

> This organization uses artificial intelligence tools to help employees work more effectively. These tools do not make decisions for the organization, and they do not reduce or transfer the accountability that employees and elected officials carry for official actions.
>
> Every employee who uses an AI tool in the course of their work remains fully responsible for the work product that results. This responsibility cannot be transferred to a software system, and it does not diminish because AI was used in the process.
>
> This principle applies to every employee, every department, and every use of AI in this organization's operations.

COMPONENT 2: DEPARTMENTAL APPLICATION

PRIMARY AUDIENCE: DEPARTMENT HEADS

What This Means for Your Department

Your employees may use approved AI tools to draft documents, research questions, prepare summaries, and support other routine work. Using those tools does not change what happens if the work product

contains an error, misrepresents a policy, or leads to a decision that later requires review. The employee who submitted the work is accountable for it.

Practically, this means two things. First, employees must review AI-generated content before they submit it as their own work — not glance at it, but actually read it, check it, and satisfy themselves that it is accurate and appropriate. Second, supervisors should set this expectation clearly and model it. If employees observe that supervisors submit AI output without review, they will follow that example.

This standard applies equally to contractors and consultants working under your direction. If an outside firm uses AI tools to produce deliverables for this organization, they are subject to the same accountability expectations.

COMPONENT 3: POLICY STANDARDS

PRIMARY AUDIENCE: CENTRAL LEADERSHIP

Policy Standards

1.1 AI as an Assistive Tool
Generative AI tools are productivity and support tools. They do not constitute an independent authority within this organization. Their outputs are not official determinations, do not carry the weight of professional judgment, and do not relieve employees of their professional obligations under law, regulation, or organizational policy.

1.2 Non-Transferable Employee Accountability
Each employee who uses AI tools in official work remains accountable for the accuracy, legality, and appropriateness of the resulting work product. This accountability cannot be delegated to an AI system. An employee may not cite reliance on AI-generated content as a defense for errors, omissions, or policy violations in official work.

1.3 AI-Generated Content as Draft Material
All AI-generated content is treated as draft material until a qualified employee reviews it, verifies it, and accepts it as accurate and appropriate. The act of submitting or relying on AI-generated content constitutes the employee's professional acceptance of that content. Review must be substantive; reviewing for formatting or obvious errors while leaving substantive content unchecked does not constitute adequate review.

1.4 Contractor and Vendor Coverage
Contractors, consultants, and vendors performing work on behalf of this organization are subject to equivalent accountability standards when using AI tools to produce organizational deliverables. Contracts for professional services should address expectations for AI use. The organization retains the same accountability for contractor deliverables as for employee work products, regardless of how those deliverables were produced.

1.5 Documentation of AI-Assisted Work
When AI tools are used in producing significant work products, employees should note AI use in the normal workflow documentation for that matter. A separate AI-specific form is not required in most circumstances. Department heads determine what constitutes significant use in their operational context. Higher-risk uses are subject to additional documentation requirements under Pillars 2 and 5.

1.6 Pillar Ownership
Primary owner: [Administrator]. Responsibility for communicating this pillar's expectations to all employees rests with department heads under the Administrator's coordination.

Legal Consideration
This pillar's contractor coverage provisions have implications for professional services agreements and vendor contracts. Consult legal counsel when drafting or revising contracts that involve AI-assisted deliverables, and when determining whether existing contracts adequately address accountability for AI-generated work product.

PILLAR 2: Prohibited Uses and Non-Delegable Decisions

COMPONENT 1: CORE PRINCIPLE

PRIMARY AUDIENCE: GOVERNING BODY

Some government decisions are too consequential to be made by software. This policy identifies those decisions and requires that qualified employees make them directly, through their own analysis and judgment.

Artificial intelligence tools may support these decisions by gathering information, drafting documents, or organizing data. They may not replace the human evaluation that determines the outcome.

The public's trust in its government depends on elected officials and their employees remaining accountable for decisions that affect people's rights, employment, safety, and access to public services.

COMPONENT 2: DEPARTMENTAL APPLICATION

PRIMARY AUDIENCE: DEPARTMENT HEADS

What This Means for Your Department

Certain decisions require your direct judgment, and AI tools cannot substitute for that judgment regardless of how capable or convenient they appear. AI may support your work in these areas by organizing information, summarizing records, or preparing drafts. The evaluation and the final determination remain yours.

The examples below illustrate where this boundary applies across common local government operations. The policy standards section that follows contains the authoritative list of restricted decision categories.

Department	AI may assist with...	AI may not determine...
Public Works / Highway / Infrastructure	Summarizing maintenance requests, drafting project scope descriptions, organizing cost data, and preparing routine inspection templates	Infrastructure project prioritization based on safety or risk; contractor performance findings; regulatory compliance determinations
Police / Sheriff / Fire / EMS	Summarizing incident data for public reports, drafting routine public communications, compiling call-for-service statistics for planning	Investigative findings; enforcement or charging recommendations; use-of-force reviews; personnel fitness determinations; findings affecting an individual's rights or liberty

Department	AI may assist with...	AI may not determine...
City / County Clerk	Drafting meeting agendas, preparing public notice templates, organizing and indexing document archives	What enters the official record; how records requests are interpreted and fulfilled; certification of official documents; election administration determinations
Health / Human Services / Veterans	Drafting outreach communications, summarizing program eligibility criteria for staff reference, and organizing case documentation templates	Individual eligibility determinations for benefits or services; clinical or diagnostic conclusions; case dispositions affecting access to public programs
Planning / Zoning / Conservation / Land Records	Summarizing application materials, drafting staff report templates, organizing permit history, and preparing public hearing notice language	Permit approvals or denials; code violation findings; variance or exception determinations; any regulatory or enforcement action affecting property rights
Central Administration (HR, Finance, Legal, Administrator)	Drafting job postings and interview questions, preparing budget narrative templates, and summarizing vendor proposals for staff review	Hiring, promotion, discipline, or termination decisions; contract award determinations; legal conclusions or formal legal opinions; expenditure authorizations requiring official approval

If you are uncertain whether a proposed AI use crosses into prohibited territory, stop and escalate before proceeding. The AI Governance Lead is your point of contact for borderline determinations. Acting first and asking later is not acceptable when the question involves a non-delegable decision.

COMPONENT 3: POLICY STANDARDS

PRIMARY AUDIENCE: CENTRAL LEADERSHIP

Policy Standards

2.1 Prohibited Decision Categories

The following categories of decisions require direct human evaluation and determination. AI tools may assist with research, drafting, or information organization related to these topics, but the substantive evaluation and final determination must be performed by a qualified employee. AI-generated content may not serve as the basis for a final decision in any of these categories without documented independent human review.

- **Personnel discipline:** Any formal corrective action, performance improvement determination, or disciplinary proceeding affecting an employee's standing, compensation, or continued employment.
- **Hiring, promotion, and termination:** Any employment decision, including initial selection, advancement, reclassification, or separation, whether voluntary or involuntary.
- **Legal conclusions:** Any formal legal interpretation, advice, or opinion presented as authoritative guidance by or on behalf of the organization.
- **Eligibility determinations:** Any determination of an individual's eligibility for public benefits, programs, permits, licenses, or government services.

- **Civil rights and equity determinations:** Any decision or formal finding involving accommodation requests, discrimination complaints, or compliance obligations under applicable civil rights requirements.
- **Law enforcement actions:** Investigative findings, enforcement priorities, use-of-force reviews, charging recommendations, or any other determination with consequences for an individual's rights or liberty.
- **Procurement and contract awards:** Bid scoring, proposal evaluation, vendor selection, and contract award determinations subject to competitive procurement requirements.
- **Regulatory and code enforcement:** Formal findings of code violations, permit non-compliance, or regulatory violations, and the penalties or remedies associated with them.
- **Closed session matters:** Any action, finding, or recommendation arising from a closed session of the governing body, including personnel matters, litigation strategy, and real estate negotiations.

> **Legal Consideration**
> The categories listed above have legal dimensions that vary depending on the type of government entity, applicable collective bargaining agreements, and state and federal law. Before finalizing or amending this list, consult legal counsel to ensure the categories are appropriately tailored to this organization's specific legal obligations and defensible under applicable law.

2.2 Permitted AI Assistance in Restricted Categories

AI tools may assist with preparatory and organizational tasks connected to each category above, including: summarizing background information, drafting template documents for employee review, organizing factual data, and identifying relevant policies or precedents for staff consideration. Permitted assistance does not alter the requirement that the substantive determination be made by a qualified human employee.

2.3 Disclosure Requirement

When AI tools are used in any preparatory capacity related to a prohibited decision category, the employee responsible for the final determination must document that they conducted an independent review and that the final determination reflects their own professional judgment. Documentation should be incorporated into the existing workflow record for the matter (personnel file, case file, contract file, etc.) and does not require a separate AI-specific form.

> **Legal Consideration**
> The disclosure requirement has potential implications for discovery in litigation, public records requests, and audit. Consult legal counsel and the records officer before finalizing how AI use documentation is integrated into existing record-keeping systems.

2.4 Equity and Bias Awareness

AI tools can produce outputs that reflect patterns in their design that treat individuals differently based on protected characteristics such as race, national origin, sex, age, or disability. Employees must be alert to outputs that appear to reflect such patterns and must not rely on those outputs in employment, service delivery, or other contexts where differential treatment of individuals is a legal or ethical concern. Before AI tools are used in any context with potential for disparate impact, the relevant department head should consult HR or the Administrator to evaluate whether the proposed use is appropriate.

> **Legal Consideration**

AI use in employment and service delivery contexts may carry exposure under applicable civil rights requirements. Consult legal counsel before AI tools are used in any high-risk context involving protected characteristics, and ask for advice on compliance obligations specific to this organization's operations.

2.5 Borderline Escalation Procedure

Employees or supervisors uncertain whether a proposed AI use falls within a prohibited category must escalate to the AI Governance Lead before proceeding. The AI Governance Lead will provide a written determination within [X] business days. Escalation determinations are recorded and reviewed periodically to identify patterns that may indicate the need for policy clarification or additional training.

2.6 Pillar Ownership

Primary owner: HR Director, in coordination with legal counsel. The HR Director is responsible for communicating prohibited use categories to all employees, maintaining the escalation log, and recommending policy revisions. The prohibited category list in Section 2.1 may not be amended without review by legal counsel and approval by the governing body.

Note to the Governing Body

By adopting this policy, the governing body establishes the categories of decision that may not be delegated to AI systems on behalf of this organization. Any proposed revision to those categories requires governing body approval. The governing body should be notified promptly if a violation of this pillar results in a personnel action, a legal claim, or a significant public records request.

PILLAR 3: Approved Environment and Tool Governance

COMPONENT 1: CORE PRINCIPLE

PRIMARY AUDIENCE: GOVERNING BODY

Not all AI tools are appropriate for government use. This policy requires that employees use only reviewed and authorized AI tools when conducting official government business, and it establishes a process for evaluating and approving those tools before they are deployed.

Consumer AI tools available to the general public do not provide the data protections, contractual safeguards, or administrative controls that government operations require. Using those tools for official work risks exposing organizational information in ways the organization cannot control or recover from.

This pillar ensures that the tools employees use for government work meet standards appropriate for public-sector operations and that the organization retains ownership and oversight of the work produced with them.

COMPONENT 2: DEPARTMENTAL APPLICATION

PRIMARY AUDIENCE: DEPARTMENT HEADS

What This Means for Your Department

The requirement is straightforward: employees may not use personal AI accounts or consumer-grade AI applications for official government work. This includes tools accessed through personal email addresses, personal subscriptions, or web-based tools with no organizational oversight or data protections. The

reason is not arbitrary — information entered into a consumer AI tool may be used by the vendor to train future AI systems, shared with third parties, or retained in ways the organization cannot control.

A critical area of attention for department heads is AI that is embedded in software your department already uses. Productivity suites, document management platforms, and other applications increasingly include AI features. Those features are not automatically approved simply because the underlying software is approved. If a tool your department uses has introduced an AI feature, notify the AI Governance Lead before staff begin using it.

If an employee in your department requests access to a new AI tool, direct them to the formal approval process. Do not authorize individual use of unapproved tools on a trial basis or under the assumption that similar tools have been approved.

COMPONENT 3: POLICY STANDARDS

PRIMARY AUDIENCE: CENTRAL LEADERSHIP

Policy Standards

3.1 Authorized Tools and Platforms
Official AI use is limited to tools that appear on the organization's authorized tool list, maintained by the [IT Director / AI Governance Lead]. Authorization is tool-specific. Approval of one AI platform does not extend to other platforms, related products, or updated versions that have not been separately reviewed.

3.2 Personal Accounts and Unapproved Applications
Employees may not enter organizational information into AI tools accessed through personal accounts, personal subscriptions, or web-based consumer applications. This prohibition applies regardless of whether the employee is working from a government device or a personal device. The controlling factor is the information, not the device.

3.3 AI Embedded in Existing Software
AI features embedded within previously approved software platforms require separate evaluation before organizational use is authorized. The default assumption is that embedded AI features are not authorized until the [IT Director / AI Governance Lead] completes an evaluation and issues an authorization or denial. Department heads are responsible for notifying the AI Governance Lead when vendors introduce AI features into platforms their departments use.

3.4 Tool Approval Process
Requests for new AI tool authorization are submitted to the [IT Director / AI Governance Lead]. Evaluation criteria include: data ownership and confidentiality protections, vendor contractual standards, security controls, administrative oversight and audit capabilities, and training data opt-out provisions. Provisional approvals may be issued for limited pilot use under defined conditions. All approval decisions are documented and communicated to affected staff. Authorization decisions are documented in writing, maintained by the [IT Director / AI Governance Lead], and communicated to affected departments before use begins.

3.5 Vendor and Contract Standards
Contracts for AI-enabled software should address: organizational ownership of data entered into the system, restrictions on vendor use of organizational data for training purposes, audit rights, incident notification requirements, and data deletion upon contract termination. Existing contracts should be reviewed for these provisions at the next renewal. Procurement of new AI-enabled software requires IT review before purchase authorization.

Legal Consideration
Vendor contracts for AI-enabled software raise data ownership, liability, and compliance questions that vary depending on contract terms and applicable law. Consult legal counsel before entering into contracts for AI tools or renewing contracts that introduce or expand AI features, particularly where the contract terms affect how organizational or constituent data is handled by the vendor.

3.6 Unauthorized Use Response

Discovery of unauthorized AI tool use — including use of personal accounts or unapproved tools for official work — should be reported to the AI Governance Lead. Response is governed by existing personnel policies. Unauthorized use resulting in actual or potential data exposure triggers incident response procedures under Pillar 8.

3.7 Pillar Ownership

Primary owner: IT Director, in coordination with the AI Governance Lead. The IT Director is responsible for maintaining the authorized tool list, conducting tool evaluations, and advising on vendor contract provisions.

PILLAR 4: Data Handling and Confidentiality Standards

COMPONENT 1: CORE PRINCIPLE

PRIMARY AUDIENCE: GOVERNING BODY

The information this organization handles belongs to the public or to the individuals it serves. Much of it is sensitive: personnel records, medical information, criminal justice data, confidential legal communications, and the private details of residents who interact with government services.

Not all of that information is appropriate to enter into AI tools. This policy establishes clear standards for what may and may not be used when interacting with AI systems, protecting both the people whose information this organization holds and the organization's legal and ethical obligations.

Most data exposure through AI systems is not the result of malicious intent. It results from employees not recognizing that the information they are working with requires protection. This pillar is primarily about awareness and clear guidance.

COMPONENT 2: DEPARTMENTAL APPLICATION

PRIMARY AUDIENCE: DEPARTMENT HEADS

What This Means for Your Department

The types of sensitive information your department handles will vary based on your function. The table below identifies examples of the categories of sensitive information common to each major department area. Employees in your department should be trained to recognize these categories and to understand that they may not enter this type of information into AI tools without explicit authorization through an approved platform with appropriate protections.

Department	Examples of sensitive information handled by this department
Public Works / Highway / Infrastructure	Infrastructure vulnerability details, contractor financial and bidding information, property owner contact information, internal project risk assessments
Police / Sheriff / Fire / EMS	Criminal justice information, arrest records, EMS and patient health data, personnel files, informant and witness information, active investigative materials, use-of-force records
City / County Clerk	Personnel records, closed session materials, personally identifiable information from public submissions, election records, confidential legal correspondence
Health / Human Services / Veterans	Medical and health information, benefit eligibility and financial data, mental health records, veterans service and disability records, child welfare case information
Planning / Zoning / Conservation / Land Records	Property owner personally identifiable information, environmental compliance records, confidential pre-application consultation materials, active enforcement case files
Central Administration (HR, Finance, Legal, Administrator)	Employee personnel and medical accommodation records, confidential legal communications and active litigation files, financial account information, closed session materials

When in doubt about whether information is sensitive, employees should assume it is and consult their supervisor or the AI Governance Lead before entering it into any AI tool.

COMPONENT 3: POLICY STANDARDS

PRIMARY AUDIENCE: CENTRAL LEADERSHIP

Policy Standards

4.1 Prohibited Data Categories

The following categories of information may not be entered into AI tools under any circumstances, including through approved platforms, without specific written authorization from the [Administrator / AI Governance Lead]:

- **Personnel and employment records:** Any document or information contained in or derived from an employee's personnel file, including performance records, discipline records, accommodation requests, and compensation information.
- **Medical and health information:** Any individually identifiable health or medical information, including information received in connection with benefit administration, accommodation requests, EMS operations, or health department services.
- **Criminal justice information:** Arrest records, investigative materials, offender records, informant or witness information, and any other information subject to criminal justice confidentiality protections.
- **Closed session materials:** Any document, communication, or record that was prepared for or arose from a closed session of the governing body.
- **Confidential legal communications:** Attorney-client communications, litigation strategy materials, and any other information subject to legal privilege.

- **Active investigative and enforcement case materials:** Records related to pending enforcement actions, open investigations, or matters that have not been publicly resolved.
- **Benefit eligibility and assistance information:** Financial, personal, or eligibility information collected in connection with applications for public benefits or assistance programs.
- **Personally identifiable information (PII):** Names, addresses, identification numbers, dates of birth, and other direct or indirect identifiers that can be linked to a specific individual, when the purpose of AI use does not require that information.
- **Information subject to a confidentiality agreement or court order:** Any information the organization is contractually or legally obligated to maintain in confidence.

> **Legal Consideration**
> Several categories above are subject to specific legal protections under federal and state law that impose independent confidentiality obligations and may carry penalties for unauthorized disclosure. Consult legal counsel to identify the specific requirements applicable to this organization's data before finalizing data handling standards and training materials.

4.2 De-identification Before AI Input

Where AI use would otherwise involve sensitive information, employees should attempt to de-identify the information before input. Adequate de-identification removes names, identification numbers, dates of birth, addresses, and other direct or indirect identifiers. When an employee is uncertain whether information has been adequately de-identified, they should consult their supervisor or the AI Governance Lead before proceeding.

4.3 Prompt Awareness

The instructions employees enter into AI tools may be stored by the platform, potentially for extended periods. Prompts containing sensitive information should be treated with the same care as the information itself. Employees should review their approved platform's documentation to understand how prompts are handled and should avoid including unnecessary sensitive details in prompts, even when using an approved tool.

4.4 Third-Party and Partner Data

Information received from other government agencies, partner organizations, or external parties that carries its own confidentiality obligations must be protected under those obligations when working with AI tools. Sharing such information with an AI platform requires the same evaluation as sharing it with any external party and must comply with any agreement or legal requirement that governs the organization's use of that information.

4.5 Pillar Ownership

Primary owner: [IT Director / HR Director / Administrator] — designated based on organizational structure. Responsibility for training employees on data handling standards rests with the HR Director and department heads.

> **Legal Consideration**
> Certain data categories handled by local governments — including health information and records connected to specific federal programs — may trigger notification, remediation, or reporting obligations if they are inadvertently entered into an AI system. Consult legal counsel to identify the specific obligations applicable to this organization before finalizing incident response procedures under Pillar 8.

PILLAR 5: Verification and Human Review Requirements

COMPONENT 1: CORE PRINCIPLE

PRIMARY AUDIENCE: GOVERNING BODY

AI tools can produce responses that sound authoritative but are factually wrong. They may cite sources that do not exist, misstate legal requirements, or calculate numbers incorrectly. This is not a malfunction — it is a known characteristic of how these systems work. The term used for this behavior is "hallucination."

This policy requires employees to review and verify AI-generated content before relying on it in official work. No AI output should be submitted, published, or acted upon as though it were automatically correct simply because a computer produced it.

The reliability of this organization's work depends on employees treating verification as a required step, not an optional one.

COMPONENT 2: DEPARTMENTAL APPLICATION

PRIMARY AUDIENCE: DEPARTMENT HEADS

What This Means for Your Department

The practical implications vary by the type of work, but the principle is consistent: before an employee uses AI-generated content in any official capacity, they must satisfy themselves that what the AI produced is accurate. This means checking facts, verifying citations, confirming that any figures or statistics are correct, and reading the content with enough engagement to catch errors that would be apparent to a qualified reader.

Verification matters most where errors matter most. A public communication with a wrong date is an inconvenience. A permit decision based on a misquoted ordinance, a budget recommendation built on an incorrect calculation, or a personnel action document that cites the wrong policy standard is a different category of problem. The examples below illustrate what verification looks like in practice across department types.

Verification Examples by Department

- **Public Works / Highway:** An employee uses AI to draft a response to a resident complaint that references a state road standard. The employee looks up the actual standard to confirm it exists and is quoted correctly before sending.
- **Police / Sheriff:** A supervisor uses AI to draft a summary of department call-for-service data for a board report. The supervisor checks the figures against the actual records before the report is submitted.
- **Clerk:** An employee uses AI to draft a public notice. The employee verifies that the legal requirements for the notice — publication deadlines, required content — match what applicable ordinance and law actually require.

- **Health / Human Services:** A staff member uses AI to prepare a summary of program eligibility criteria for internal reference. The summary is checked against the actual program guidelines before it is distributed to staff.
- **Planning / Zoning:** A planner uses AI to draft a staff report. All ordinance section references, land use classifications, and required findings are verified against the actual ordinance before the report is finalized.
- **Central Administration:** HR uses AI to draft a job posting. The classification requirements, salary range, and required qualifications are verified against the official position description and compensation plan before posting.

Supervisors should understand that modeling verification is as important as requiring it. If employees observe that supervisors accept AI output without checking it, no policy language will change that behavior.

COMPONENT 3: POLICY STANDARDS

PRIMARY AUDIENCE: CENTRAL LEADERSHIP

Policy Standards

5.1 Verification as a Required Workflow Step

Verification of AI-generated content is a required step before the content is submitted, published, presented, or otherwise relied upon in official work. Verification is not optional and may not be waived for convenience or time pressure.

5.2 Risk-Based Review Standards

The depth of verification required corresponds to the potential consequences of error. Three tiers apply:

- **Routine use:** Internal administrative tasks with no external disclosure, no sensitive data, and no regulatory consequence. Examples: drafting internal meeting agendas, organizing internal reference summaries. User review required before use.
- **Standard use:** External communications, formal internal reports, staff recommendations, policy research, and other work that will be seen by others or inform decisions. Examples: public-facing communications, grant research summaries, staff briefings. Thorough user review required; supervisor awareness expected for significant work products.
- **High-risk use:** Content touching legal obligations, financial figures, eligibility or benefit determinations, regulatory requirements, or any work product likely to become part of the public record or be cited in an official proceeding. Examples: formal recommendations to the governing body, budget documents, compliance determinations. Independent verification of all specific claims required; supervisor or subject-matter review before reliance.

5.3 Specific Verification Requirements

Regardless of risk tier, employees must independently verify any AI-generated content that includes:

- Factual claims that will be represented as accurate in official communications or decisions
- Legal citations, statutory references, ordinance provisions, or regulatory quotations
- Numerical figures, calculations, cost estimates, or financial projections
- Policy or ordinance interpretations

- References to external standards, studies, guidance documents, or court decisions
- Names, titles, dates, or other specific identifying information

5.4 Supervisory Review Threshold
Supervisors must review AI-assisted work products before submission or reliance when: the work product will be presented to the governing body; the work product will be filed in an official regulatory or legal proceeding; or the department head or supervisor independently determines that the work product carries elevated risk based on its content or intended use.

5.5 Verification Documentation
For high-risk uses, employees should document that they completed independent verification as part of the normal workflow record. A brief notation in the file is sufficient. Separate AI-specific documentation is not required unless the work product will be subject to formal audit or legal review.

5.6 Pillar Ownership
Primary owner: Department Heads, supported by the [Administrator / AI Governance Lead]. Department heads are responsible for ensuring verification expectations are understood and followed within their teams.

> **Legal Consideration**
> Reliance on AI-generated content in official proceedings, formal filings, legal documents, or public records without adequate verification may expose the organization to liability for errors in those materials. The standard of care applicable to professional verification varies by role and context. Consult legal counsel when establishing verification standards for high-risk uses specific to this organization's operations and legal obligations.

PILLAR 6: Records, Retention, and Disclosure Governance

COMPONENT 1: CORE PRINCIPLE

PRIMARY AUDIENCE: GOVERNING BODY

The work this government does is public work. When employees use AI tools in the course of their official duties, the materials they produce — and in many cases the instructions they gave the AI tool — may be subject to the same public records obligations that govern any other government document.

This policy requires that AI-assisted work be identified, retained, and managed consistently with the organization's existing records obligations. The method of producing a document does not change the organization's responsibility to manage it appropriately.

Residents who interact with this government have a right to access public records. That right is not diminished by the fact that AI played a role in creating those records.

COMPONENT 2: DEPARTMENTAL APPLICATION

PRIMARY AUDIENCE: DEPARTMENT HEADS

What This Means for Your Department

The practical standard is this: if you would retain a document or record it in your department's files without AI involvement, you retain and record it with AI involvement. The fact that AI helped produce the content does not change what the document is or what your obligations are with respect to it.

There are three specific situations your department should be prepared for. First, if your department receives a public records request, any AI-assisted documents responsive to that request must be identified and provided like any other record. Second, if an AI tool helped produce instructions, analysis, or recommendations that were used in an official decision, those materials may be part of the record for that decision. Third, if a matter your department is involved in becomes the subject of a legal claim, complaint, or investigation, notify the Administrator and legal counsel immediately — AI-related materials connected to that matter may be subject to preservation requirements.

Questions about whether specific AI-generated materials are subject to records obligations should go to the [Records Officer / Clerk] before a records request arrives, not after.

COMPONENT 3: POLICY STANDARDS

PRIMARY AUDIENCE: CENTRAL LEADERSHIP

Policy Standards

6.1 AI-Assisted Work Products as Records
AI-assisted work products that are created or received in connection with the transaction of official government business are subject to this organization's records management obligations. The fact that content was generated with AI assistance does not alter its records status. Employees should apply the same records management judgment to AI-assisted work products as they apply to all official work.

6.2 Prompts as Records
Prompts entered into AI tools in connection with official work may constitute records depending on their content and connection to official business. Employees should treat prompts that direct substantive official work with the same care as the resulting work product. The default operating posture is retention: prompts connected to official work should be treated as records unless the [Records Officer / Clerk] determines otherwise. Questions about specific situations should be directed to the [Records Officer / Clerk].

> **Legal Consideration**
> The application of public records law to AI-generated content, including prompts and outputs, is an evolving area. Current guidance at the state level continues to develop. Consult legal counsel to understand how this organization's records obligations apply to AI-assisted work before finalizing records management procedures.

6.3 Retention Schedule Application
Existing retention schedules apply to AI-assisted work products based on the subject matter and function of the work, not the method of production. The [Records Officer / Clerk] is responsible for advising departments on how existing schedules apply to AI-assisted materials and for recommending any amendments needed to address new record types that do not clearly fall within existing categories.

6.4 Public Records Requests
AI-assisted documents responsive to a public records request must be identified and produced consistently with the organization's standard records request procedures. Employees may not withhold or fail to identify responsive records on the basis that they were AI-generated or AI-assisted. Questions about

the scope of disclosure for specific AI-related materials must be directed to the [Records Officer / Clerk] and legal counsel before the response deadline.

6.5 Litigation Holds and Investigations

When a matter involving AI-assisted work becomes the subject of litigation, a formal complaint, or an official investigation, the responsible employee and supervisor must notify the [Administrator / legal counsel] immediately. All applicable AI-related materials — including prompts, outputs, and related communications — are subject to preservation obligations consistent with the organization's litigation hold procedures and must not be deleted or modified pending resolution.

> **Legal Consideration**
> Litigation hold obligations apply to all records, including AI-generated materials, from the point at which litigation is reasonably anticipated — not only after a lawsuit is filed. The scope and duration of those obligations depend on the nature of the claim and applicable law. Consult legal counsel at the first indication that a matter may result in litigation or a formal proceeding.

6.6 Official Documents and Legal Instruments

AI tools may assist in drafting ordinances, resolutions, contracts, and other official legal instruments. No official legal instrument may be finalized or executed based solely on AI-generated content without review and approval by qualified staff and, where applicable, legal counsel. The use of AI in drafting formal legal documents raises questions specific to each document type; legal counsel should be consulted before AI tools are used in this context.

6.7 Pillar Ownership

Primary owner: [Records Officer / City or County Clerk], in coordination with legal counsel. The Records Officer or Clerk is responsible for advising on retention obligations, responding to records requests involving AI-related materials, and recommending retention schedule amendments.

> **Note to the Governing Body**
> The governing body should understand that AI-assisted work products produced in connection with official government business may be subject to public records requests, and that the organization's records obligations are not diminished by the use of AI tools in producing those records. The governing body should be notified if a public records request or legal proceeding surfaces AI-related materials that raise significant policy or legal questions requiring governing body awareness.

PILLAR 7: Workforce Readiness and Access

COMPONENT 1: CORE PRINCIPLE

PRIMARY AUDIENCE: GOVERNING BODY

Adopting a policy is not the same as building the capacity to follow it. This pillar ensures that employees receive the training and organizational support they need to use AI tools responsibly before they begin using them for official work.

Access to AI tools is granted based on training completion, not simply on interest or availability. Employees who have not completed required training do not have access to AI tools for official use.

This pillar also addresses the organization's obligations to its workforce in connection with AI adoption — including the obligation to communicate transparently with employees and to comply with

applicable collective bargaining requirements before implementing AI tools that affect working conditions.

COMPONENT 2: DEPARTMENTAL APPLICATION

PRIMARY AUDIENCE: DEPARTMENT HEADS

What This Means for Your Department

Your primary responsibility under this pillar is ensuring that no employee in your department uses AI tools for official work before completing required training. Access to approved tools is contingent on training completion, and you are responsible for knowing which of your employees have met that requirement.

Your supervisory staff needs training that goes beyond what frontline employees receive. Supervisors must understand not only how to use AI tools but also how to monitor use, enforce verification standards, recognize problems, and support employees who have questions or concerns during implementation.

If your department includes represented employees and AI tools will affect how work is done, do not proceed with implementation before consulting with the Administrator and HR about collective bargaining obligations. The sequence matters: engage the required process first, then implement.

When your department introduces AI tools or expands AI use, communicate proactively with your team about what is changing, what is expected, and who to contact with questions. Changes introduced without explanation generate resistance and misuse.

COMPONENT 3: POLICY STANDARDS

PRIMARY AUDIENCE: CENTRAL LEADERSHIP

Policy Standards

7.1 Baseline Training Requirements
All employees authorized to use AI tools for official work must complete baseline training before access is granted. Baseline training covers: how generative AI works and its limitations (including hallucination); approved tools and how to access them; data handling and confidentiality standards (Pillar 4); verification requirements (Pillar 5); prohibited uses (Pillar 2); and the process for escalating questions or reporting concerns. Training completion is documented and records maintained by [HR Director / Administrator]. Refresher training is required at [X]-year intervals or when policy changes materially.

7.2 Tiered Access Framework
Initial access may be limited to lower-risk applications while employees develop familiarity with AI tools and organizational expectations. Expanded access to higher-risk or more capable AI tools requires completion of role-appropriate training. Access levels are determined by [AI Governance Lead / IT Director] in consultation with department heads. Access may be suspended or revoked for violation of this policy or for demonstrated inability to use AI tools within policy standards.

7.3 Supervisory Training
Department heads and supervisors require training that addresses their additional responsibilities: monitoring employee AI use, enforcing verification standards, managing escalations, and supporting teams through implementation. Supervisory training is a prerequisite for department-level AI deployment. Supervisors who have not completed supervisory-level training may not authorize AI use within their teams.

7.4 Elected Officials and Appointed Board Members

Elected officials and appointed board members who use AI tools in connection with their official duties are subject to the same data handling and confidentiality standards applicable to employees. Orientation materials calibrated to governing body responsibilities should be made available. Governing body members cannot be required to complete training as a condition of service but should be offered training and encouraged to participate.

7.5 Competency and Acknowledgment

Upon completing baseline training, employees acknowledge in writing that they have read and understand this policy and their responsibilities under it. Acknowledgment is incorporated into the organization's standard policy acknowledgment process and is renewed when this policy is materially revised.

7.6 Implementation and Change Management

Deployment of AI tools within any department should be preceded by communication to affected employees identifying: what tools are being introduced, what official uses are authorized, what training is required, and who to contact with questions or concerns. Employees should have a designated point of contact during implementation. Feedback mechanisms should be established so employees can raise concerns without fear of reprisal.

7.7 Workforce and Labor Considerations

AI implementation that affects working conditions, workload standards, performance evaluation criteria, or employee monitoring may implicate collective bargaining obligations for represented employees. Administration should consult with [HR Director / legal counsel] before implementing AI tools in any context that may affect the terms and conditions of employment for represented employees. Transparent communication with all employees — represented and non-represented — about AI adoption plans, intended uses, and policy expectations is essential to effective and accountable implementation.

> **Legal Consideration**
> AI implementation in workplaces with collective bargaining agreements may raise mandatory bargaining obligations depending on how the technology affects employees' terms and conditions of employment. The scope of those obligations depends on applicable collective bargaining agreements and applicable labor law. Consult legal counsel and engage labor relations advisors before implementing AI tools in any department with represented employees, and before those implementations are communicated to staff.

7.8 Pillar Ownership

Primary owner: HR Director, in coordination with the [Administrator / AI Governance Lead]. The HR Director is responsible for developing and maintaining training materials, documenting training completion, and advising on workforce and labor considerations.

PILLAR 8: Governance, Oversight, and Continuous Improvement

COMPONENT 1: CORE PRINCIPLE

PRIMARY AUDIENCE: GOVERNING BODY

A policy that is adopted and then left unattended does not protect the organization. This pillar establishes the ongoing governance structure that keeps this framework active and effective: who is

responsible for administering it, how problems are identified and addressed, how the policy is updated as technology and law evolve, and what the governing body can expect to hear about the organization's AI activities.

Governance is not a one-time action. It is a continuing organizational responsibility. The adoption of this policy is the beginning of that responsibility, not the end of it.

COMPONENT 2: DEPARTMENTAL APPLICATION

PRIMARY AUDIENCE: DEPARTMENT HEADS

What This Means for Your Department

Your primary obligations under this pillar are reporting and awareness. If something goes wrong with AI use in your department — an employee uses a prohibited tool, relies on inaccurate AI output in official work, enters sensitive data into an AI system, or uses AI to make a non-delegable decision — you are responsible for reporting it to the AI Governance Lead promptly and accurately. The organization cannot improve what it does not know about.

You also play a role in making policy governance work over time. Your operational experience with AI tools is valuable input to the policy review process. When the AI Governance Lead or administration solicits feedback on how this policy is working in practice, provide candid responses. Policies that do not reflect operational reality do not get followed.

COMPONENT 3: POLICY STANDARDS

PRIMARY AUDIENCE: CENTRAL LEADERSHIP

Policy Standards

8.1 AI Governance Lead

The organization designates [Title] as the AI Governance Lead. This individual is responsible for: maintaining the authorized tool list; managing the tool approval process; receiving, logging, and responding to escalation requests; coordinating employee training; tracking and reporting incidents; and recommending policy revisions to the Administrator. The AI Governance Lead reports to [Administrator / designated official]. For smaller organizations, the AI Governance Lead function may be combined with existing IT or HR leadership responsibilities.

8.2 AI Governance Committee

Organizations with multiple departments or significant AI use should establish an AI Governance Committee to provide cross-functional oversight. Recommended membership: [Administrator or designee], IT Director, HR Director, [Records Officer / Clerk], legal counsel, and at least one department head representative on a rotating basis. The Committee meets at least [X] times per year and reviews: incident reports, tool approval requests, policy compliance, training effectiveness, and policy revision recommendations. For smaller organizations, governance committee functions may be performed by the existing leadership team with AI governance added as a standing agenda item.

8.3 Supervisory Oversight Function

Department heads and supervisors are responsible for monitoring AI use within their teams for compliance with this policy. Supervisors should be alert to: employees using unapproved tools, submitting AI-generated content without adequate review, entering sensitive data into AI systems, or using AI in

connection with non-delegable decisions. Supervisors are not expected to technically audit AI usage; they are expected to observe work product quality and to take appropriate action when concerns arise.

8.4 Incident Response and Reporting

An AI-related incident includes any of the following:

- Entry of prohibited or sensitive data into an AI tool
- Reliance on materially inaccurate AI output in an official work product
- Use of an unapproved AI tool for official government work
- Use of AI in connection with a prohibited decision category under Pillar 2
- Any AI-related event that results in or may result in harm to an individual, legal exposure for the organization, or significant public concern

Reporting procedure: Employees must report suspected incidents to their supervisor and the AI Governance Lead promptly upon discovery. The AI Governance Lead logs all incidents and determines the appropriate organizational response. Incidents involving potential legal liability, data exposure, or harm to individuals are escalated to the Administrator and legal counsel immediately. All incident records are reviewed periodically to identify policy gaps, training deficiencies, and patterns requiring organizational response.

> **Legal Consideration**
> Incident response obligations vary depending on the type of incident, the data involved, and applicable law. Certain incidents involving specific categories of protected information may trigger mandatory notification or remediation requirements with defined timelines. Consult legal counsel to identify the specific obligations applicable to this organization before finalizing incident response procedures, and establish a protocol for rapid legal consultation when incidents occur.

8.5 Governing Body Reporting

The [Administrator / AI Governance Lead] provides the governing body with an annual report on the organization's AI governance activities. The annual report includes at minimum: a summary of authorized tools in use, employee training completion status, incident summary (with identifying details removed), escalation request summary, and recommended policy revisions. The governing body is notified promptly — outside the annual cycle — when any incident triggers legal counsel involvement, a formal complaint, significant public scrutiny, or a state or federal inquiry.

8.6 Policy Review Triggers

Scheduled review: This policy is reviewed at least every [X] months by the AI Governance Lead and [Administrator], with findings presented to the governing body.

Event-driven review is required promptly when any of the following occur:

- A significant AI-related incident within the organization
- New guidance from state or federal authorities with implications for local government AI governance
- A significant change in the AI tools used by the organization or a material change in vendor terms
- A court decision or enforcement action relevant to AI use in government operations
- A new collective bargaining cycle in which AI-related working conditions are subject to negotiation
- A governing body request for policy review

8.7 Performance Indicators

The AI Governance Lead tracks and reports the following indicators to assess policy effectiveness and identify areas requiring attention:

- Employee training completion rate, by department
- Number and category of incidents reported
- Number and disposition of escalation requests
- Number and outcome of tool approval requests
- Any public records requests or legal proceedings involving AI-related materials

8.8 Pillar Ownership

Primary owner: [Administrator], with operational coordination by the AI Governance Lead. The Administrator is ultimately responsible for ensuring this policy is implemented, that incidents are addressed appropriately, and that the governing body receives the reporting it needs to exercise its oversight role.

> **Note to the Governing Body**
> The governing body is the ultimate governance authority for this policy. It adopts the policy, approves material revisions, and receives annual reports on implementation. The governing body should expect to be notified promptly of any AI-related incident that triggers legal counsel involvement, a formal complaint, or significant public concern. Governing body members with questions about AI governance activities between annual reports should direct them to the [Administrator].