

Resolution No. 26-233

A resolution adopting the City of Arlington's Artificial Intelligence Policy

WHEREAS, the State of Texas has enacted new requirements governing the use of artificial intelligence by local governments, including Texas Government Code Chapter 2054, Subchapter S, the Texas Responsible Artificial Intelligence Governance Act, and Texas Department of Information Resources rules under 1 TAC Chapter 219; and

WHEREAS, the City desires to comply with state requirements; NOW THEREFORE

BE IT RESOLVED BY THE CITY COUNCIL OF THE CITY OF ARLINGTON, TEXAS:

I.

That the City Council hereby adopts the City of Arlington's Artificial Intelligence Policy establishing governance requirements for the responsible use of artificial intelligence systems by City departments, employees, and vendors.

II.


A substantial copy of the City of Arlington's Artificial Intelligence Policy is attached hereto and incorporated herein for all intents and purposes.

PRESENTED AND PASSED on this the 23rd day of June, 2026, by a vote of 9 ayes and 0 nays at a regular meeting of the City Council of the City of Arlington, Texas.


ATTEST:


MARTHA GARCIA, City Secretary




JIM R. ROSS, Mayor

APPROVED AS TO FORM:
MOLLY SHORTALL, City Attorney

BY 



Artificial Intelligence Policy

VERSION 1.0

Artificial Intelligence Policy

Authorized By: City Council
Scope: Citywide
Review Frequency: As needed
Effective Date: Upon approval
Supersedes: N/A (First Issue)

I. Document and Information History

Publication Version History

Version numbering: Pre-publication drafts are 0.x. The first published version is 1.0. Minor revisions increment the decimal (e.g., 1.1). Major content upgrades increment the whole number (e.g., 2.0).

Version	Date	Author	Description
1.0	June 6, 2026	Information Technology, Office of Cybersecurity	Initial publication

II. Purpose

The Artificial Intelligence Policy (“Policy”) provides a framework for the thoughtful, secure, and innovative use of artificial intelligence systems, including generative AI technologies, across City of Arlington (“City”) operations. It is designed to help the City explore and adopt AI in ways that enhance public services, improve accessibility, increase efficiency, and strengthen resident engagement, while protecting sensitive information and aligning with applicable Texas law.

III. Scope

All City officials, employees, departments, vendors, contractors, volunteers, and third-party service providers who develop, deploy, operate, or interact with artificial intelligence systems on behalf of the City are subject to this Policy.

This Policy applies to all artificial intelligence systems including, but not limited to:

- Generative AI systems
- Machine Learning Models
- Automated Decision Systems
- Agentic AI Systems
- Predictive Analytics Tools
- Computer Vision Systems
- Natural Language Processing Tools
- Any software or system meeting the definition of an artificial intelligence system under Texas Government Code §2054.003(1-a)

IV. Effective Date and Implementation

This citywide policy becomes effective on the date it is approved by City Council.

V. Roles and Responsibilities

- **Department Directors** – Responsible for ensuring that employees within their department comply with the requirements of this Policy.
- **Information Technology Department / Office of Cybersecurity (OCS)** – Responsible for oversight of AI security, privacy review, policy interpretation, and incident response coordination. The City Chief Information Security Officer, or their designee, shall serve as the City's AI Risk Officer.
- **City Manager** – Shall approve any deployment of any Heightened Scrutiny AI system. May adopt more granular standards that fall within the scope of this Policy.

VI. Definitions

The following terms carry specific meaning under Texas law. Definitions are provided to help the reader understand how each term applies to this Policy. Complete statutory definitions are maintained in the City's Information Security Policy Glossary.

- **Artificial Intelligence System** – Any machine-based system that, for any explicit or implicit objective, infers from the inputs the system receives how to generate outputs, including content, decisions, predictions, or recommendations, that can influence physical or virtual environments.
- **Consequential Decision** – A decision that has a material legal or similarly significant effect on the provision, denial, or conditions of a person's access to a government service.
- **Controlling Factor** – A factor that is: (A) the principal basis for making a consequential decision; or (B) capable of altering the outcome of a consequential decision.
- **Heightened Scrutiny Artificial Intelligence System** – An AI system specifically intended to autonomously make, or be a controlling factor in making, a consequential decision. Does not include systems intended to perform a narrow procedural task; improve the result of a previously completed human activity; perform a preparatory task to an assessment relevant to a consequential decision; or detect decision-making patterns or deviations.
- **Principal Basis** – The use of an output from a Heightened Scrutiny AI system to make a decision without: (A) human review, oversight, involvement, or intervention; or (B) meaningful consideration by a human.
- **Unlawful Harm** – Any condition in which the use of an AI system results in a consequential decision that causes harm to an individual who is a member of a state or federally protected class in violation of law.

VII. Policy

7.1 Human Review of AI Outputs

- Employees using AI systems shall verify the accuracy and appropriateness of AI-generated outputs before relying on them in any official City capacity, with a level of review commensurate to the risk and impact of the output.
- For Heightened Scrutiny AI systems, the review must constitute meaningful human evaluation of outputs and decisions. These systems require increased human oversight relative to lower-risk systems.
- AI systems must be capable of being disabled until harmful or inaccurate decision-making can be remedied.

7.2 AI Risk Classification and Inventory

- AI systems shall be classified according to their potential impact on residents, City services, and decision-making processes.
- AI system outputs that are the principal basis for, or a controlling factor in, consequential decisions shall be classified as Heightened Scrutiny AI systems and require a written risk assessment and additional governance review prior to deployment.
- The City shall maintain a current inventory of AI systems deployed within City operations, including the purpose of each system, the department responsible for its use, its risk classification, and the oversight and mitigation measures in place.
- The City shall identify and inventory all Heightened Scrutiny AI systems at least annually and whenever a new AI system is deployed or an existing system is materially modified.
 - The City shall maintain the inventory in accordance with applicable technical standards established for this purpose.
- Upon request by the Texas Department of Information Resources, the City shall provide the City's review of the deployment and use of Heightened Scrutiny AI systems in the manner DIR prescribes.

7.3 Written Risk Assessment for Heightened Scrutiny AI Systems

a. Before the City develops, procures, deploys, or uses a Heightened Scrutiny AI system, and when a material change is made to the system, a written risk assessment shall be conducted to evaluate the probability and severity of harm that could result from implementation.

b. The written risk assessment shall consider and document:

- The system's known security risks and available mitigation steps;
- Performance metrics, including measurements of accuracy and relevance of outputs, and operational aspects such as model latency, uptime, and error rate;
- Transparency, including information about the system's algorithms and decision-making processes, the data used to train the model, and the availability of inputs and outputs to monitor decision-making over time;
- The potential for the system to result in unlawful harm, including the risk of discriminatory outcomes.

c. The City Manager shall review the completed risk assessment prior to deployment and provide a written decision to approve or deny the deployment.

d. The City shall maintain a record of the completed written risk assessment and all relevant documents for as long as required by the applicable records retention schedule.

7.4 Acceptable Use Policies for Heightened Scrutiny AI Systems

- Prior to deployment of any Heightened Scrutiny AI system, the deploying department shall develop and adopt an acceptable use policy specific to that system, identifying approved use cases, system limitations, and prohibited uses.
- All employees and contractors with access to the Heightened Scrutiny AI system must be trained on the acceptable use policy prior to system use.

7.5 Data Security and Privacy

Use of AI systems shall be consistent with all City security, technology, data governance, and records management policies, procedures, and standards.

- Sensitive or personally identifiable information shall not be entered into AI systems that are publicly accessible or controlled by third parties unless approved through Information Technology Department review.
- Employees shall be trained on the risk of inputting sensitive or personally identifiable information into publicly available AI systems that may use inputs to train the model or share inputs with other users outside the City. All training must be delivered through a DIR-certified artificial intelligence training program.
- The City must understand what personally identifiable information an AI system uses, how that information has been and will be collected, and how the system uses, stores, and shares such information with third parties, prior to using any City-held personally identifiable information in the system.
- The City must collect and maintain only that personally identifiable information needed for operations and must establish a process to delete such information consistent with records retention schedules and other legal requirements.

7.6 Prohibited Uses of AI

Pursuant to Texas Business and Commerce Code Chapter 552 (Texas Responsible AI Governance Act), the following restrictions apply to the use of City AI systems:

- Deceptive AI identity: the City shall not represent AI systems as human when interacting with the public, unless a reasonable person would already know the interaction involves an AI system.
- Social scoring: the City shall not use AI to evaluate or classify a person or group based on social behavior or personal characteristics in a way that results in detrimental treatment unrelated to, or disproportionate to, the observed behavior.
- Biometric identification: the City shall not develop or deploy an AI system for the purpose of uniquely identifying a specific individual using biometric data, or through targeted or untargeted gathering of images or other media, without the individual's consent, where doing so would infringe on any constitutional or legal right.
- Manipulation of human behavior: the City shall not develop or deploy an AI system that intentionally aims to incite or encourage a person to commit self-harm, harm another person, or engage in criminal activity.
- Unlawful discrimination: the City shall not develop or deploy an AI system with the intent to unlawfully discriminate against any protected class.

7.7 Resident Disclosure Duties

- When the City deploys or uses an AI system that is public-facing or that is a controlling factor in a consequential decision, the City shall include a standardized notice on all related applications, websites, and public computer systems.

- The City shall use the standardized notice form developed by the Texas Department of Information Resources to disclose that artificial intelligence is being used.
- The standardized notice shall be included on all related applications, internet websites, and public computer systems associated with the AI system.

7.8 AI Transparency and Documentation

- AI systems deployed by the City that are public-facing or that are a controlling factor in a consequential decision shall include a standardized notice as required by Texas law.
- When AI systems are used in consequential decisions, the responsible City department shall maintain sufficient documentation to allow City staff to understand the system's role in the decision-making process.
- The City shall collaborate with developers of AI systems and require transparency to understand how a system operates, the source of the data the system was trained on, and its intended use cases.
- The City shall not represent AI systems as human when interacting with the public.
- The City shall maintain records of the sources of data used in each AI system and how the system is modified throughout its lifecycle.

7.9 Public Records Compliance

- Records generated, used, or stored by AI systems shall be managed in accordance with the Texas Public Information Act and applicable records retention requirements.
- Employees using AI systems must maintain or be able to retrieve records of inputs, prompts, and outputs consistent with applicable records retention schedules and the Texas Public Information Act.
- The responsible City department shall establish appropriate retention schedules for AI system records and consider the Public Information Act implications related to the storage of data inputs and outputs.

7.10 Employee AI Training

- City employees who use a computer to complete twenty-five percent (25%) or more of their required job duties shall complete a DIR-certified artificial intelligence training program on an annual basis.
- The City shall identify qualifying employees and coordinate enrollment in a DIR-certified AI training program on an annual basis.
- Department Directors are responsible for ensuring employees within their department complete required training within the timeframe established by Human Resources.
- Departments deploying Heightened Scrutiny AI systems shall provide system-specific training to all employees and contractors who access, use, or manage the system, covering identified risks and appropriate mitigation methods.

7.11 Security and Incident Reporting

- Any breach, unauthorized access, or misuse of an AI system, including entry of confidential data into public or uncontrolled AI systems, shall be reported to the Information Technology Department in accordance with the City's Cybersecurity Incident Response Plan.
- City employees who identify concerns regarding the accuracy, fairness, or potential harm of a Heightened Scrutiny AI system's output shall report those concerns to their Department Director and to the Information Technology Department.

- Incidents involving potential violations of the prohibited uses identified in Section 7.6 shall be escalated to the City Attorney's Office.

7.12 Resident Remedies

- The City shall provide a mechanism for individuals to seek remedy when an AI system makes a consequential decision that unfairly impacts them in a material way.
- Each department deploying an AI system shall designate a point of contact for individuals seeking information about a consequential decision.
- Departments shall develop internal procedures to identify and remedy negative impacts caused by the use of AI systems.

7.13 Fairness and Ongoing Evaluation

- When conducting the written risk assessment for a Heightened Scrutiny AI system, the deploying department shall consider the potential for the system to result in unlawful harm, including the risk of discriminatory outcomes, as part of the probability and severity analysis.
- The City shall ensure its use of AI systems does not infringe upon the legally protected rights and liberties of the individuals it serves or result in unlawful harm.
- The City shall implement data governance practices for AI systems throughout the AI system's lifecycle to ensure fairness.
- Deployed AI systems shall be evaluated on an ongoing basis to ensure continued benefit and to identify concerns regarding accuracy, fairness, or harm. Evaluations shall be documented.

7.14 AI Procurement Review

- Prior to entering a contract for the acquisition, licensing, or deployment of any Heightened Scrutiny AI system, the deploying department shall submit the proposed system to the Office of Cybersecurity for review.
- The AI Risk Officer shall evaluate the proposed system for alignment with this Policy, applicable Texas law, and City security, privacy, and data governance standards prior to contract execution.
- Heightened Scrutiny AI systems shall require completion of the written risk assessment and notification of the City Manager before procurement may proceed.
- Vendor contracts for AI systems shall include provisions requiring compliance with applicable Texas AI requirements, including the AI Code of Ethics and, for Heightened Scrutiny AI systems, implementation of the NIST AI Risk Management Framework.
- Vendor contracts shall provide the City with at least thirty (30) days' notice and opportunity to cure for any identified non-compliance with applicable AI requirements.
- The City shall: (1) comply with all procurement requirements established by the Department of Information Resources for information resources and technology that include a Heightened Scrutiny AI system; (2) consider conducting an impact assessment aligned with state agency requirements when deploying a Heightened Scrutiny AI system; and (3) review relevant resources posted by DIR on its website.

7.15 Accountability

While the City may delegate tasks and decision-making to AI systems, City staff remain accountable for the decisions AI systems make and the outcomes they produce.

- The City shall ensure its vendors are contractually bound to the AI Code of Ethics and any relevant laws or regulations governing the use of AI systems.

- The City shall ensure AI systems it deploys comply with the legal obligations the City has at both the state and federal level.

VIII. Exceptions

All exception requests shall be submitted to the Information Technology Department for review and initial classification of the risk impact level. If applicable, the City Attorney's Office may also review the request in accordance with regulatory requirements.

Risk Level	Actions Required	Risk Register
Low	Update policy language to prevent similar exceptions.	No
Medium	Written acceptance required from the responsible Department Director.	Yes
High	Approval from the City Manager's Office.	Yes

IX. Non-Compliance

For City staff, non-compliance with citywide policies and standards is subject to disciplinary action in accordance with the Human Resources Personnel Policy Manual.

For vendors, non-compliance may result in suspension of access, termination of contract(s), legal action, and loss of future business with the City. Vendor contracts involving artificial intelligence systems must comply with applicable Texas AI requirements, including but not limited to Texas Government Code Chapter 2054, Subchapter S, and 1 TAC Chapter 219.

X. Approval Authority

City Council holds approval authority for this Policy. The City Manager, or their delegate, may adopt and implement more granular standards that fall within the scope of this Policy.

The Information Technology Department is responsible for policy maintenance and implementation.

XI. Table of References

- Texas Government Code Chapter 2054, Subchapter S – Artificial Intelligence Governance
- Texas Business and Commerce Code Chapter 552 – Texas Responsible AI Governance Act
- 1 TAC Chapter 219 – Texas DIR Rules on AI Use by Governmental Entities