



**TOWN OF LOS GATOS**  
**COUNCIL POLICY COMMITTEE REPORT**

MEETING DATE: 08/25/2026

ITEM NO: 2

---

DATE: August 25, 2026  
TO: Council Policy Committee  
FROM: Chris Constantin, Town Manager  
SUBJECT: Review and Provide Direction on a Proposed Artificial Intelligence Policy Framework

**RECOMMENDATION:**

Review and provide feedback on the proposed Artificial Intelligence Policy Framework and recommend forwarding the Framework to the Town Council for consideration and adoption.

**BACKGROUND:**

Artificial intelligence is increasingly being incorporated into many aspects of software, research tools, communications platforms, analytical systems, and operational technologies used by public agencies. In many cities such as Los Gatos, the Microsoft 365 software suite includes AI capabilities for certain tasks. Use of other tools, including free versions of commercial tools, are in use throughout many organizations without proper policies, training, and guidance.

These tools can improve service delivery, expand staff capacity, increase accessibility, and reduce routine work. They also can create issues if not properly guided and managed. To reduce exposure in our organization, staff are working on creating the appropriate structure to leverage the benefits of AI tools and capabilities, and to moderate and place boundaries around how those tools and capabilities are used.

Staff reviewed examples of practices from other cities as well as government and public sector guidance from the GovAI Coalition and National Institute of Standards and Technology to understand how other organizations, such as Los Gatos, could implement these types of tools while considering how to protect from misuse. AI use in government is an emerging field, and most cities tend to be underdeveloped in how the tools should be properly used. Thus, while we look at other practices, the approach for Los Gatos should be Los Gatos-focused and tailored to how we want to leverage these new tools.

**PREPARED BY:** Chris Constantin, Town Manager and Sai Kim, IT Director

---

Reviewed by: Town Manager and Town Attorney

---

SUBJECT: Review and Provide Direction on the Proposed Artificial Intelligence Policy Framework

DATE: August 25, 2026

Staff developed the proposed Artificial Intelligence Policy Framework as a concise Council policy that establishes the Town's governing principles and boundaries for AI use. The Framework is intentionally technology neutral and does not establish detailed operating requirements. The Framework is a balance between being too sparse and overly broad in language to allow administrative policies and procedures to expand upon areas of the Policy Framework. Implementation would be addressed through administrative policies, procedures, and standards established by the Town Manager and supported by the IT Director. Staff training on the Framework will be required, and a review of the tools the Town would authorize would consider the requirements of the Framework.

Operational use of public safety technologies would continue to be governed by applicable law and Police Department policies and procedures, together with generally applicable Town requirements. The proposed Framework is included as Attachment 1, and supporting practices from other organizations are included in Attachment 2.

#### DISCUSSION:

##### *Framework Approach*

The Framework is organized around seven policy pillars that substantially reflect the principal themes found in the municipal policies and other governance materials reviewed by staff. The pillars include

- Public Value and Workforce Augmentation,
- Human Accountability and Oversight,
- Privacy and Data Stewardship,
- Security and Reliability,
- Transparency and Trust,
- Equity and Non-Discrimination, and
- Legal Compliance.

The Framework also establishes the respective governance roles of the Town Council, Town Manager, Town Attorney, and Information Technology Director and provides direction and boundaries for internal administrative policies and Townwide AI use. The Framework also defines Artificial Intelligence and AI-Enabled Technology broadly enough to include AI functionality embedded within otherwise ordinary products, defines the threshold for a Consequential Town Decision or Action, establishes four prohibited use categories, and requires the Town Manager to review the Framework at least every two years and recommend amendments as needed.

The central policy choice is to encourage responsible AI use while keeping consequential governmental authority with Town employees. AI may assist with analysis, detection, alerts,

monitoring, data collection, recommendations, and decision support, but those functions do not themselves constitute a final Town decision or action. When AI contributes to a consequential Town decision or action, the decision or action remains subject to the human review, confirmation, and authorization required by applicable Town or departmental policy.

The Framework also uses a proportionate disclosure standard so routine AI assistance does not require disclosure when an employee reviews the work, exercises independent judgment, and accepts responsibility for the result. The disclosure standard reflects a deliberate policy choice. The City and County of San Francisco 2025 Generative AI Guidelines use a comparable tiered approach. Disclosure is not required for internal drafting or support. For public facing or sensitive work, the Guidelines require additional documentation, and individuals impacted by AI systems must receive direct notice when Generative AI substantially contributes to the work product. Other reviewed approaches are broader. The 2023 Seattle Generative AI Policy required attribution when AI generated text was used substantively in a final product and for AI generated images and video, Los Altos Hills requires attribution whenever generated content is published, and Santa Cruz asks employees to indicate when AI contributed substantially to a work product. Staff recommends proportionate standards and intends to provide more guidance in administrative policies and procedures which balance the extent of AI use with how the resulting output is used, interpreted, or relied upon by decision-makers and the public. This allows staff to have a more thoughtful discussion and to provide staff a workable standard that can be implemented.

#### *Prohibited Uses of AI*

The prohibited use section is intentionally technology neutral and establishes four broad boundaries. The four areas address clear boundaries that are considered in other policies and industry guidance and are framed to be more applicable to our local government environment.

AI may not independently make or carry out a consequential Town decision or action. A human must remain responsible for reviewing and authorizing such decisions or actions, a concept commonly described as keeping a human in the loop. Town staff should also confirm that information materially relied upon is sufficiently accurate and complete for the intended use.

AI may not be used in a manner that results in unlawful discrimination, or that is deceptive, coercive, or intended to improperly manipulate an individual. While this may seem obvious, it is important to establish a clear restriction that can be implemented through administrative policies, procedures, and guidelines designed to reduce the risk of these practices.

AI may not be used to identify, monitor, classify, or profile individuals through automatic and continuous means unless the use is consistent with applicable law and authorized by applicable Town or departmental policy. Portland expressly restricts surveillance or profiling outside lawful and approved purposes, while Santa Cruz addresses related risks through restrictions on

unauthorized access, facial recognition, social scoring, and certain automated decisions. For Los Gatos, the intent is to preserve lawful and authorized uses while establishing a clear policy boundary for uses that could affect privacy and civil liberties.

AI may not be used to access, use, or disclose protected Town information except when authorized for an appropriate Town purpose and consistent with applicable law and Town or departmental policy. These provisions address the underlying risks reflected in municipal policies such as San José, Portland, and Santa Cruz while allowing more specific technology and use restrictions to be established through Town administrative policies, Police Department policies and procedures, or other departmental requirements consistent with the Framework. Local government information is subject to public records, privacy, confidentiality, security, and other legal requirements. Some Town information is protected from disclosure or restricted in use, so it should only be used with AI in a manner authorized for the information and the Town purpose involved.

*Comparison with AI Policies, Practices, and Guidance*

Staff reviewed the publicly available municipal AI policies and governance materials<sup>1</sup>. Although the documents vary in level of detail and organizational placement, the principal policy themes have many consistencies with the proposed Los Gatos Framework. It should be noted that many of the cities and organizations that have AI policies are larger or intended for larger organizations. In addition, many cities do not have AI policies publicly available, indicating that many municipalities have not yet formally addressed AI governance through a Council-adopted policy. By adopting the proposed Framework, the Town of Los Gatos would take a proactive approach to AI governance and be among the municipalities beginning to establish a formal framework for the responsible use of AI.

**City of San José - Artificial Intelligence Policy 1.7.12:** San José's policy is organized around effectiveness, transparency, equity, accountability, human-centered design, privacy, security and safety, and workforce empowerment. The Los Gatos Framework addresses these core governance subjects through its seven pillars and governance structure, although San José expressly includes workforce education and training within Workforce Empowerment. San José also requires AI reviews, vendor technical documentation, incident response, and sunset procedures. San José identifies specific prohibited uses, including real-time or covert biometric identification, emotion analysis, fully automated consequential decisions without meaningful human oversight, social scoring, and cognitive behavioral manipulation. Los Gatos addresses the underlying concerns through broader prohibited use boundaries and its Human Accountability and Oversight pillar rather than adopting the same technology specific list.

**City of Sacramento – Artificial Intelligence Policy:** Sacramento's citywide policy, effective January 2026, uses substantially the same core principles as San José and includes formal AI

---

<sup>1</sup> Some of the best practices summaries utilized AI-assisted tools.

SUBJECT: Review and Provide Direction on the Proposed Artificial Intelligence Policy Framework

DATE: August 25, 2026

review, human review of generative AI content, training, and specific privacy requirements for public safety employees. Los Gatos addresses the same underlying Council level concerns through employee accountability, review appropriate to the consequences of error, protection of sensitive information, and human confirmation of consequential actions. Los Gatos also more explicitly preserves Police Department operational policies while maintaining generally applicable Town requirements.

**City of Anaheim – Administrative Regulation 158:** Anaheim treats AI as a resource to support, rather than replace, human judgment and accountability. Its regulation also addresses privacy and security, verification of AI output, transparency, bias, and department specific rules, and separately requires employee acknowledgment. The substantive governance concepts are reflected in the Los Gatos pillars, definitions, and prohibited uses. Employee acknowledgment, tool approval, and detailed data handling are administrative implementation measures that Los Gatos intends to address through administrative and technical standards.

**City of Santa Cruz - Artificial Intelligence Appropriate Use Policy:** Santa Cruz emphasizes beneficial use, privacy and security, human oversight, accuracy, transparency, accountability, equity, continuous learning, and limits on fully automated consequential decisions. Los Gatos addresses the same underlying concerns through broader requirements concerning human accountability, lawful and nondiscriminatory use, authorized identification and monitoring, and protection of Town information. Santa Cruz also includes more specific requirements concerning public interaction, tool approval, employment uses, and IT risk review; those subjects can be addressed in the Town's administrative AI Use Policy and departmental procedures.

**Town of Los Altos Hills – Generative Artificial Intelligence Acceptable Use Policy:** Los Altos Hills adopted a Council policy focused specifically on generative AI in February 2024. It encourages innovation while requiring human responsibility, fact-checking, public records compliance, training, and attribution. Los Gatos incorporates the same underlying accountability, reliability, transparency, legal compliance, and workforce principles, but is broader and technology neutral so that it remains applicable as AI becomes embedded in ordinary software and operational systems.

**City of Portland - BTS 4.04 Artificial Intelligence Use and Governance:** Portland's 2026 administrative rule identifies human-centered use, security and safety, privacy, transparency, equity, accountability, effectiveness, and workforce empowerment and applies a risk-based review model. Portland also identifies specific prohibited uses addressing unlawful activity, unauthorized surveillance or profiling, automated consequential decisions without appropriate human review, manipulation, social scoring, and deceptive content. These risks are also covered by Los Gatos's broader prohibited use boundaries. Portland's procurement, vendor documentation, risk-model, risk assessment, and operational oversight requirements, together

with its use of inventories or summaries of approved uses as transparency mechanisms, illustrate the type of implementation detail that the Los Gatos Framework would delegate to administrative policy.

**City of Boston - Generative Artificial Intelligence Policy:** Boston places ultimate responsibility for AI-assisted work on employees, requires training, identifies unacceptable uses, and supplements its policy with separate resources including use-case guidance, an FAQ, an AI inventory, and data security guidance. Los Gatos takes the same general approach of supporting productive use while retaining human accountability and using separate administrative controls for implementation.

*Alignment with Broader Government and Public Sector Practices*

The proposed Framework is also aligned with the broader guidance reviewed by staff. NIST AI Risk Management Framework 1.0 emphasizes valid and reliable systems, safety, security and resilience, accountability and transparency, privacy, fairness, and risk-based governance. The GovAI Coalition template uses substantially the same principles and contemplates supporting implementation tools and procedures, including AI review, vendor technical documentation, incident response, and an AI Governance Manual. California Department of Technology guidance likewise addresses risk assessment, privacy, security, procurement, inventory, acceptable use, and workforce training. These materials support the Town's approach of establishing durable Council principles while reserving operational controls for administrative implementation.

California Department of Technology Letter 25-01 establishes Generative AI requirements for State entities and is used here as a public sector benchmark rather than as a Town mandate. The Framework's Legal Compliance pillar requires compliance with applicable law, and the Town Attorney will advise on legal requirements as implementation proceeds. Implementing administrative policies and procedures should include verification of compliance with applicable State and federal law.

CONCLUSION:

The proposed Artificial Intelligence Policy Framework provides a concise and durable Council-level foundation for the Town's use of AI. The Framework incorporates the principal policy elements found in the municipal policies and broader best-practice materials reviewed by staff. Staff does not recommend adding additional operational detail to the Council Framework and recommends that the Council Policy Committee review the Framework, provide any desired feedback, and recommend forwarding it to the Town Council for consideration and adoption by Resolution.

PAGE 7 OF 7

SUBJECT: Review and Provide Direction on the Proposed Artificial Intelligence Policy Framework

DATE: August 25, 2026

COORDINATION:

The framework was reviewed by Town department heads and amendments were made based on their input. Prior to final Town Council consideration, staff will continue coordination with the Town Attorney, Information Technology, Human Resources, and Police Department regarding implementation of the Framework. Staff will confirm with the Town Attorney the Town's obligations under applicable California artificial intelligence legislation and under the Meyers-Milias-Brown Act, including any notice or meet and confer obligations arising from the introduction of AI-enabled technologies, and will provide notice to recognized employee organizations as appropriate.

Attachments:

1. Artificial Intelligence Policy Framework
2. AI Practices and Policy Reference Compilation