



Town of Chester AI Governance Policy

Last Updated: 5/13/2026

Approved: 7/27/2026

SECTION 1: Purpose & Scope

This policy establishes the rules, responsibilities, and guardrails for the use of artificial intelligence (AI) tools by anyone acting on behalf of the Town of Chester, Massachusetts. Its purpose is to enable the responsible adoption of AI while protecting resident data, preserving public trust, and ensuring compliance with applicable law.

1.1 Why This Policy Exists

AI tools are increasingly embedded in the software and services the Town uses every day. Some of these tools are obvious — a chatbot you type into. Others are invisible — AI features running inside your email client, permitting software, or HR platform without your awareness. This policy exists to ensure that every use of AI by or on behalf of the Town is intentional, documented, and governed.

1.2 Who This Policy Applies To

This policy applies to all of the following individuals when acting in their capacity on behalf of the Town of Chester:

All full-time and part-time Town employees

Elected officials, including Select Board members

Appointed committee and commission members

Volunteers performing work on behalf of the Town

Contractors, consultants, and temporary workers with access to Town systems or data

Any third-party vendor whose AI tools process Town data

Throughout this policy, the term "covered individuals" refers to all persons listed above.

1.3 What This Policy Governs

This policy covers the selection, approval, use, and oversight of all AI systems used in connection with Town business, including but not limited to:

Generative AI tools (chatbots, writing assistants, image generators)

AI features embedded in existing software (permitting, email, HR, finance)
Meeting transcription and summarization tools
Agentic AI tools that take actions on behalf of users
No-code/low-code automation tools with AI capabilities
AI used to inform, support, or make decisions affecting residents, employees, or operations

1.4 Relationship to Other Policies and Laws

This policy operates alongside and does not replace the following:

Town of Chester Acceptable Use Policy
Town of Chester Records Retention Schedule (per Secretary of the Commonwealth guidelines)
Massachusetts Public Records Law (M.G.L. ch. 66, § 10)
Massachusetts Open Meeting Law (M.G.L. ch. 30A, §§ 18–25)
Massachusetts Wiretap Statute (M.G.L. ch. 272, § 99) — all-party consent
Massachusetts Data Breach Notification Law (M.G.L. ch. 93H)
Massachusetts Consumer Protection Act (M.G.L. ch. 93A) as applied to AI per the Attorney General's April 2024 Advisory
Any applicable federal regulations (e.g., ADA, FERPA, HIPAA) depending on the data or population involved

⚠ MASSACHUSETTS AG ADVISORY ON AI (APRIL 2024)

The Massachusetts Attorney General has issued formal guidance confirming that existing consumer protection (Chapter 93A), civil rights, and data privacy laws apply fully to the use of AI systems. Any use of AI that affects residents — including service delivery, benefits determinations, communications, and enforcement — must comply with these laws. This policy incorporates that guidance throughout.

SECTION 2: Definitions

Every term below has a specific, practical meaning in this policy. When in doubt, refer back here.

2.1 Types of AI Systems

Term	Definition
Copilot / Assistant AI	AI built into tools you already use that helps you draft, summarize, or search. Examples: Microsoft Copilot in Outlook or Teams, Notion AI, Google Gemini in Gmail. You control when it runs and what it does.
Standalone AI	A separate AI tool you open and interact with directly. Examples: ChatGPT, Claude, Gemini.com. You type a prompt and read the response.
Embedded / Integrated AI	AI that is woven into other software and may run automatically without you clicking anything. Examples: AI features inside your permitting software, HR platform, or accounting software that summarize records,

	flag anomalies, or draft responses on their own.
Agentic AI Tool	AI that takes actions on your behalf — not just generating text, but doing things: browsing the web, reading and writing files, sending emails, clicking buttons, or running workflows. Examples: AI browsers (like Claude in Chrome), Cowork, multi-step automation agents. These tools often run in the background and may have persistent access to your accounts. They carry the highest risk level.
Meeting Assistant	A tool that joins a meeting and records, transcribes, or summarizes it. Examples: Otter.ai, Fireflies, Fathom, AI notetakers built into Zoom or Teams. These capture everything said, including Sensitive Data shared during the call.
Sanctioned AI System	Any AI tool formally reviewed and approved by the Town of Chester. The Approved Tools list is maintained in Section 12 of this policy.
Unsanctioned / Shadow AI	Any AI tool not on the Approved list, including personal accounts of otherwise approved tools. Using your personal ChatGPT account at work is Shadow AI even if the Town uses ChatGPT Team.

2.2 What Is Ethical Output?

'Ethical output' is not a vague aspiration. It has three measurable layers in this policy.

Layer 1 — Prohibited Output Categories

AI output must not be used if it:

Contains fabricated facts, statistics, names, citations, or quotes presented as real

Reflects or amplifies bias against any person or group based on a protected characteristic

Exposes confidential, proprietary, or personally identifiable information

Is deceptive, designed to mislead about its origin, accuracy, or intent

Reproduces copyrighted material without authorization

Layer 2 — The Pre-Use Checklist

Before using any AI-generated output, confirm all five of the following:

#	Check	Pass/Fail
1	ACCURATE: I have verified key facts and figures against a reliable source.	☐
2	UNBIASED: The output does not unfairly favor or disadvantage any person or group.	☐
3	DATA-SAFE: No confidential, resident, or employee data is in this output.	☐
4	ATTRIBUTED: If shared externally, AI involvement is disclosed or confirmed as	☐

	outlined in Section 7.	
5	REVIEWED: For all AI used in outputs (financial, legal, HR, external), a human read and approved before action was taken.	☐

Layer 3 — Scoring Rubric (High-Stakes Outputs)

Use this when AI output informs a major decision, is sent to a board/committee, or involves employee or sensitive/critical data.

Dimension	What to Assess	Min Score	Your Score
Accuracy	Are all facts and figures verifiable? Have you checked all claims?	2 / 3	___ / 3
Bias	Does the content treat all people and groups fairly?	2 / 3	___ / 3
Tone	Is the tone appropriate for the audience and purpose?	2 / 3	___ / 3
Attribution	Is it clear what was AI-generated? Has the source been noted internally?	2 / 3	___ / 3
Data Safety	Does the output contain data that should not have been in this AI system? (PII, employee or resident data, financial data)	3 / 3	___ / 3

SCORING GUIDE

1 = Does not meet standard (do not use) | 2 = Meets standard with minor reservations | 3 = Fully meets standard

If any dimension scores 1, the output must be revised or discarded. A Data Safety score below 3 requires immediate escalation to the Town Administrator or the Select Board.

2.3 Key Terms Glossary

Term	Definition
Traceability	The ability to reconstruct the full story of an AI-assisted decision: what tool, what input, what output, who reviewed, and what action was taken. If you cannot answer all five, the decision is not traceable.
Transparency (Internal)	Everyone inside the Town who needs to know that AI was involved in a process is informed before they act on the output.
Transparency (External)	Residents or counterparties are informed when AI plays a material role in communication or decisions that affect them. See Section 7 for required disclosure language.

AI-Assisted Decision	Any decision where AI output was a contributing factor, even if a human made the final call. Example: A manager reads an AI summary of a performance review and uses it — that is an AI-assisted HR decision.
Sensitive Data	Resident names, employee names, contact info, financial data, health information, employee records, contracts containing financial terms, personal information, or confidential provisions, and anything marked Confidential or used solely internally.
Human-in-the-Loop	A qualified human actively reviews and approves AI output before it is used or acted upon. Forwarding or copying AI output without reading it does not count.
Persistent Access	When an AI tool maintains ongoing authorization to access your accounts, files, or systems — even when you are not actively using it. Examples: an OAuth connection, a saved API key, a browser agent with file system access. Persistent access requires explicit authorization and regular review.
Prompt	The instruction or question you give to an AI system. A prompt containing Sensitive Data means that data has left your control and may be stored or processed by a third party.
Autonomous Action	An action taken by an agentic AI tool without a human clicking 'confirm' for each step. Examples: an agent that sends an email on your behalf, or browses a website and fills out a form.
Covered Individual	Any person subject to this policy. See Section 1.2 for the complete list.

SECTION 3: Approved & Prohibited Uses

Use the matrix below to determine whether your intended AI use is approved. When unsure, default to 'Requires Review' and contact the Town Administrator or the Town's IT Provider.

ALWAYS OK	REQUIRES REVIEW	NEVER OK
Drafting internal emails or memos using no Sensitive Data	Drafting resident-facing proposals, contracts, or recommendations	Entering resident, employee, or financial data into a personal or unsanctioned AI account
Summarizing public information, news, or general research	Summarizing internal reports or documents that contain Town data	Recording a meeting without informing and receiving consent from all participants first (see M.G.L. ch. 272, § 99)
Brainstorming and idea generation for human review	Using AI to screen, rank, or score candidates or employee performance	Sending AI-generated content externally without human review
Writing or reviewing notes when no Sensitive Data is included	Publishing AI-generated content under the Town's name	Using AI to automatically execute financial, legal, or HR decisions
Using AI copilots for personal	Enabling a new AI feature inside an	Granting an agentic tool persistent

productivity on non-sensitive tasks	existing approved tool (needs review first)	access without documented approval
Using an approved agentic tool on pre-approved, non-sensitive workflows	Using an agentic tool on any workflow that touches Sensitive Data	Allowing an agentic tool to send communications or take financial actions without a human confirmation step
	Using AI to draft social media posts, press releases, or newsletter content (see Section 7.4)	Using AI to make or materially inform decisions about residents (permit approvals, benefits eligibility, code enforcement, service routing) without a documented civil rights and bias review
		Using AI to draft emergency or public safety communications without real-time oversight and sign-off by the responsible official (see Section 7.4)

WHEN IN DOUBT

Treat any uncertain use case as 'Requires Review.' Contact the Town Administrator and the Town's IT Provider before proceeding.

There is no penalty for asking. There may be consequences for proceeding without approval.

SECTION 4: Shadow AI & Unsanctioned Tools

Now that you know what is and is not approved, this section addresses what to do if you are already using something that is not. Shadow AI is the most common and most overlooked AI risk for municipalities.

4.1 What Counts as Unsanctioned?

Your personal ChatGPT, Claude, Gemini, or Copilot account (even if the Town uses the enterprise or team version)

Any browser extension with AI features not on the Approved list

A meeting recorder or note-taker installed or enabled by a vendor on your behalf

AI features inside SaaS tools that were not reviewed when the parent tool was approved

An AI browser or agentic tool (e.g., Cowork, Claude in Chrome) not on the Approved list

Any tool a colleague recommended that is not on the Approved list

IMPORTANT CLARIFICATION

Using an unsanctioned tool is not automatically a disciplinary matter. The risk is that these tools may send your data to third-party servers, train on your inputs, or take actions on your accounts without your full awareness. The goal of this section is to surface that risk, not punish curiosity.

4.2 What To Do If You Are Currently Using an Unsanctioned Tool

Stop entering Sensitive Data into it immediately.

Submit it for review using the process in Section 4.3.

You may continue using the tool only for public/general data tasks for up to 10 business days while the review is pending. If the tool is Declined, cease all use immediately.

For agentic tools: revoke any active connections or OAuth authorizations until the tool is reviewed.

Expect a response within 10 business days.

4.3 Tool Approval Process

Step	Action
Step 1	Email the IT Provider and Town Administrator with the requested AI tool. Include: tool name, vendor, what data you plan to use in it, and your business justification. Answer: 'What problem will this tool solve?' and 'How will this tool improve the work I do?' Include both quantitative and qualitative metrics. For agentic tools, also describe what actions the tool will take and what systems it will connect to.
Step 2	Town Administrator, in coordination with the IT Provider, reviews the vendor's data processing agreement, privacy policy, and the Vendor Risk Assessment Checklist (see Section 11.5) within 10 business days.
Step 3	Tool is either Approved, Conditionally Approved (with restrictions), or Declined (with explanation).
Step 4	Approved tools are added to the list in Section 12. All covered individuals are notified. Any persistent access granted to agentic tools is documented and scheduled for review per Section 6.2.

SECTION 5: Data Handling Rules

This section answers: what can I put into an AI tool? The answer depends on both the type of data and the type of AI tool. The authorization rules in Sections 3 and 4 tell you which tools you may use. This section tells you what you may put into them.

CORE RULE

If data would be sensitive in an email to a stranger, it is sensitive in an AI prompt. Agentic and embedded tools raise the stakes further: they may access, copy, or act on your data automatically, even without an explicit prompt from you. See Section 6 for the full rules on agentic tools.

5.1 Data Classification & AI Input Rules

Data Type	Examples	AI Input Rule	Agentic / Embedded Tools
Public / General	Industry news, public reports, general knowledge	OK in any sanctioned tool	OK with documented workflow
Internal / General	Process docs, meeting notes (no resident names)	OK in sanctioned tools; not in personal accounts	Requires workflow review — see Section 6
Resident Data	Resident names, contacts, project details, contracts containing financial terms, personal information, or confidential provisions	NEVER in unsanctioned tools. Sanctioned only with manager approval.	NEVER without documented approval — see Section 6
Employee Data	Performance reviews, salaries, HR records	NEVER in any AI tool without explicit Town Administrator and IT Provider approval	NEVER — see Section 6
Regulated / Legal	Health info, legal advice, compliance data, financial advice	NEVER without written sign-off from Town Administrator and IT Provider	NEVER — see Section 6

5.2 Meeting Assistants — Special Rules

AI meeting transcription and summarization tools carry a specific risk that most people are not aware of: when you use an unauthorized tool, the recording, transcript, and everything discussed in that meeting — including any Sensitive Data — may be sent to the vendor's servers, used to train AI models, retained indefinitely, and potentially shared with third parties. This is not a hypothetical. It is standard practice for many free and premium AI meeting tools.

For this reason, meeting transcription tools are subject to stricter rules than other AI tools. The rules below cover three situations: when you are hosting a meeting, when you are attending a meeting hosted by someone else, and what to do when a conversation shifts unexpectedly to confidential topics.

⚠ MASSACHUSETTS WIRETAP STATUTE — M.G.L. ch. 272, § 99

Massachusetts is an ALL-PARTY CONSENT state. Recording any oral or wire communication without the consent of ALL parties is a criminal offense punishable by up to 5 years imprisonment and a \$10,000 fine. This applies to AI meeting assistants. An employee who activates an AI recording or transcription tool without obtaining explicit consent from every participant is not just violating this policy — they may be

committing a felony under Massachusetts law. There are no exceptions for 'internal only' meetings or 'just taking notes.'

⚠ OPEN MEETING LAW — M.G.L. ch. 30A, §§ 18–25

For public boards, committees, and commissions subject to the Open Meeting Law:

- AI-generated transcripts and summaries of public meetings are presumed to be public records subject to disclosure under M.G.L. ch. 66, § 10. Store and manage them accordingly. See Section 8 for full public records guidance.
- AI transcription of executive sessions creates a significant risk that privileged content could become discoverable. Do NOT use AI transcription tools in executive sessions without prior written approval from Town Counsel.
- Official minutes of public meetings must still be created and approved per Open Meeting Law requirements. An AI transcript does not substitute for official minutes unless formally adopted by the body as its minutes.
- If an AI tool is used to draft minutes of a public meeting, a designated member must review, edit, and formally approve those minutes before they are posted.

When you are hosting a meeting:

Only use a meeting transcription or summarization tool that is on the Approved Tools list (Section 12). All other tools are prohibited for meetings where Sensitive Data may be discussed.

Before starting any recording or transcription, notify all participants: 'This meeting will be recorded and transcribed by an AI tool.' Do not begin until you have received explicit consent from all participants. This is required by Massachusetts law (M.G.L. ch. 272, § 99).

If any participant declines consent, stop the recording immediately — no exceptions. Do not proceed and record anyway.

Use the waiting room feature for all virtual meetings. Deny entry to any unrecognized or external meeting assistant bot. If an external participant attempts to join with their own AI transcription tool, explain that only authorized tools may be used in meetings hosted by the Town and ask them to disable it.

Handle all transcript and summary output as Sensitive Data: store files on authorized systems, apply appropriate access controls, and share only with intended recipients.

Validate the output of the AI meeting assistant for accuracy before distributing or acting on it. Do not forward a transcript or summary without reviewing it first.

Do not use meeting transcripts as input into other AI tools without manager approval.

When you are attending a meeting hosted by an external party:

If the meeting will involve Sensitive Data and the host is using an AI transcription or summarization tool you have not reviewed, do not consent to that recording.

If the external host insists on using their AI transcription tool, offer one of two alternatives: suggest moving the meeting to a platform hosted by the Town where only approved tools

are in use, or move the sensitive portion of the conversation to a different channel (phone call, secure message, in-person).

Do not disclose Sensitive Data in a meeting where an unauthorized AI tool is actively recording, even if the host assures you the data is safe.

When a conversation shifts to confidential topics mid-meeting:

If a meeting that started as routine shifts into discussion of restricted or confidential information, and a recording or transcription is active, ask the host to pause or stop the recording before you disclose any Sensitive Data.

If the host cannot or will not stop the recording, excuse yourself from that portion of the meeting and follow up through a secure channel.

SECTION 6: Agentic & Embedded AI Tools

WHY THIS SECTION EXISTS

Most AI policies are written for chatbots: you type, AI responds, you decide what to do with the answer. Agentic tools are fundamentally different. They act. They browse the web, read your files, send emails, fill out forms, and run multi-step workflows — often without asking you to confirm each step.

Embedded AI tools carry similar risk: they are woven into software you already use, and they may be analyzing your data and taking actions in the background right now, without you realizing it.

The rules in this section exist because these tools interact with real systems, real data, and real people on your behalf. A mistake by an agentic tool is not just a wrong answer — it can be a sent email, a submitted form, or an exposed file.

6.1 Examples of Agentic & Embedded AI Tools

Tool Type	Description & Risk
AI Browsers	Browsers with built-in AI agents that can browse websites, extract information, fill out forms, or complete tasks autonomously on your behalf. Example: Claude in Chrome, Arc with AI. Risk: The agent can access any site you are logged into and take actions using your credentials.
Desktop Automation Agents	Tools that control your computer — clicking buttons, opening applications, reading files, and completing multi-step workflows. Example: Cowork. Risk: These tools have file system and application access that extends well beyond what a chatbot can reach.
No-Code Workflow Agents	Tools that connect your apps and run automated sequences triggered by events. Examples: Zapier AI, Make.com with AI steps, Copilot Studio. Risk: Misconfigured workflows can silently move, expose, or delete data across connected systems.
Embedded Permitting / HR /	AI features built into tools like your permitting software, tax collection

Finance AI	system, constituent management tool, or HR platform that automatically summarize records, draft responses, or flag anomalies. Risk: The AI may be processing Sensitive Data automatically, even if no one explicitly enabled it.
AI Email & Calendar Tools	AI that reads your inbox or calendar and takes actions — drafting replies, scheduling meetings, unsubscribing, or summarizing threads. Risk: Your inbox contains Sensitive Data. An agent with inbox access can read and act on all of it.

6.2 Rules for Agentic Tools

Before enabling any agentic tool:

The tool must be on the Approved list (Section 12). Submit a Tool Request (Section 4.3) if it is not.

Document exactly which systems the tool will connect to and what actions it is authorized to take.

Get explicit approval from the Town Administrator and IT Provider before granting access to any system that contains Sensitive Data.

Ensure the tool is configured with the minimum access required — never grant broader permissions 'just in case.'

While using an agentic tool:

For any workflow involving Sensitive Data or external communications, require a human confirmation step before the tool executes.

Do not allow an agentic tool to send emails, submit forms, make purchases, or take irreversible actions autonomously without a human-in-the-loop checkpoint.

Monitor the tool's action log regularly to confirm it is operating within its approved scope.

If the tool behaves unexpectedly or takes an action it was not supposed to, stop it immediately and notify the Town Administrator. Follow the AI Incident Response procedure in Section 9.

Persistent access:

Document every OAuth connection, API key, or saved credential granted to an agentic tool.

All persistent access authorizations for agentic and embedded AI tools must be reviewed and re-authorized or revoked every 90 days. The IT Provider is responsible for conducting this review in coordination with the Town Administrator. This is the governing rule — all other references to persistent access reviews in this policy defer to this section.

When a covered individual leaves the organization or their role changes, immediately revoke all agentic tool access tied to their accounts.

Never share an agentic tool's credentials or API keys with another person.

⚠ THE INJECTION RISK

A specific risk with agentic tools is prompt injection: malicious instructions hidden in a webpage,

document, or email that the AI agent reads and then executes as if you had given the command.

Example: An AI browser agent is sent to research a vendor. The vendor's website contains hidden text instructing the agent to forward your contact list to an external address. The agent complies.

Mitigation: Only allow agentic tools to take actions that require explicit human confirmation. Review action logs after every run. Report any unexpected behavior immediately.

6.3 Embedded AI in Existing Tools

Many software tools you already use — your permitting software, email client, timekeeping system, project management tool — have added AI features over time. Some of these are enabled by default.

Your responsibilities:

Periodically check the AI or automation settings in the tools you use. If you see an AI feature you did not intentionally enable, report it to the IT Provider.

Before using any AI feature inside an existing tool for a new purpose (e.g., using your constituent management tool's AI to draft a resident email for the first time), confirm that use case is covered by the tool's existing approval.

If an embedded AI tool has access to Sensitive Data as a result of its integration, that access must be documented and reviewed on the same 90-day schedule as other agentic tools (see Section 6.2).

SECTION 7: Traceability & Transparency Standards

The previous sections define what you may use and what you may put into AI tools. This section defines what you must be able to show afterward, and what you must disclose to others. These are the two governance practices that make AI use auditable and trustworthy.

7.1 Traceability — What It Means and When It Applies

Traceability means you can reconstruct the story of an AI-assisted decision after the fact. It is required for any 'Requires Review' use case and any use involving Sensitive Data.

A traceable AI-assisted decision answers all five of these questions:

Question	What to Record
What tool?	Name and version of the AI system used (e.g., 'ChatGPT Team, May 2026')
What input?	The prompt or data given to the AI, or a clear description of it
What output?	The AI response that was used, or a summary of it

Who reviewed?	The name of the person who reviewed the output before action was taken
What decision?	What action was taken as a result, and when

For agentic tools, traceability also requires:

A log of the actions the agent took (files accessed, emails sent, forms submitted, websites visited)

Confirmation that no action exceeded the defined scope of the approved workflow

The name of the person who authorized the agent to run

WHERE TO STORE RECORDS

A brief note in the relevant project file, permit record, email thread, or departmental log is sufficient, as long as it answers the five questions.

For financial and HR decisions, records must be retained per the Massachusetts Municipal Records Retention Schedule maintained by the Secretary of the Commonwealth's office. Consult the Town Clerk or Town Administrator if you are unsure of the applicable retention period.

7.2 Transparency — Internal

If you use AI to produce content that will be used or acted on by someone else inside the Town, indicate that AI was involved. A note in a document header, an email line, or a verbal mention is sufficient.

INTERNAL DISCLOSURE (required)

"Note: This [summary / draft / analysis] was generated with AI assistance and reviewed by [Name] on [Date]."

7.3 Transparency — External

External transparency applies when AI played a material role in something sent to or done on behalf of a resident, vendor, or business partner.

EXTERNAL DISCLOSURE (required when applicable)

"Parts of this [communication / proposal / analysis] were prepared with the assistance of AI tools and reviewed by our team for accuracy."

WHEN EXTERNAL DISCLOSURE IS NOT REQUIRED

You do not need to disclose AI involvement when:

AI was only used for grammar, spelling, or formatting and did not affect substance, OR

The output was fully rewritten by a human and AI contributed no recognizable content.
When unsure, disclose.

7.4 Social Media & Public Communications

AI-generated content used in public-facing communications carries unique risk because it reaches residents directly and represents the Town's official voice. The following rules apply:

All AI-generated social media posts, press releases, newsletter content, and website copy must be reviewed and approved by the relevant department head before publishing.

AI must never be used to draft emergency communications — including public safety alerts, weather emergencies, evacuation notices, or Amber Alerts — without real-time human oversight and sign-off by the responsible official.

AI-generated content published on Town websites, social media accounts, or distributed to mailing lists must meet WCAG 2.1 accessibility standards. This includes appropriate alt text for images, readable formatting, and plain language.

Do not use AI to generate content that could be interpreted as legal advice, medical guidance, or binding policy positions without review by Town Counsel or the appropriate authority.

If a resident or media outlet asks whether specific Town content was AI-generated, answer honestly.

SECTION 8: Public Records & Retention

LEGAL BASIS — M.G.L. ch. 66, § 10

Under the Massachusetts Public Records Law, any documentary material made or received by a government entity is presumed to be a public record unless a specific exemption applies. This includes digital records, drafts, and communications — regardless of the platform or tool used to create them.

8.1 AI Records as Public Records

AI prompts, outputs, traceability logs, and related documentation created in connection with Town business are likely public records subject to disclosure under M.G.L. ch. 66, § 10. This means:

AI-generated drafts of Town communications, memos, or reports — even if not finalized — may be subject to public records requests.

The prompts you type into an AI tool while conducting Town business may be responsive to a public records request.

Traceability records (Section 7.1) serve a dual purpose: they support good governance and they may be required for public records production.

AI-generated meeting transcripts and summaries of public meetings are presumed to be public records.

AI-generated records from executive sessions are subject to the restrictions in Section 5.2 — consult Town Counsel before producing them in response to a public records request.

8.2 Handling Public Records Requests for AI Content

When the Town receives a public records request that may involve AI-generated content:

Search all sanctioned AI tools and storage locations for responsive records, including prompts, outputs, and traceability logs.

If responsive records exist in an AI tool's conversation history, export and produce them in a standard format.

Apply any applicable exemptions (e.g., attorney-client privilege for legal consultations, personnel records exemption) — consult with Town Counsel as needed.

Respond within 10 business days as required by M.G.L. ch. 66, § 10(b).

If you are unsure whether AI-generated content is responsive, consult the Town Clerk or Town Administrator.

8.3 Retention of AI Records

AI-related records must be retained in accordance with the Massachusetts Municipal Records Retention Schedule maintained by the Secretary of the Commonwealth's office. The retention period depends on the type of record the AI output supports:

Financial records supported by AI analysis: retain per the schedule applicable to financial records.

Personnel records involving AI: retain per the schedule applicable to personnel records.

General correspondence assisted by AI: retain per the schedule applicable to correspondence.

AI traceability logs: retain for the same period as the underlying decision they document.

Do not delete AI conversation histories, prompts, or outputs that relate to Town decisions until the applicable retention period has expired.

⚠ WHY PERSONAL AI ACCOUNTS ARE A RECORDS PROBLEM

Town business conducted on personal AI accounts (your personal ChatGPT, Claude, Gemini, etc.) may still be subject to public records requests — but those records may be difficult or impossible for the Town to locate, export, or produce. This is one of the key reasons Section 4 prohibits using personal AI accounts for Town business. If you have already used a personal account for Town work, export and preserve any relevant records immediately and notify the Town Administrator.

SECTION 9: AI Incident Response

This section defines the procedure for responding to incidents where an AI tool causes or contributes to a data exposure, unauthorized action, or harm to the Town, its residents, or its employees.

9.1 Types of AI Incidents

An AI incident includes, but is not limited to:

Data exposure: Sensitive Data (resident, employee, or financial information) was entered into an unsanctioned AI tool or a sanctioned tool in a manner that violates this policy.

Unauthorized autonomous action: An agentic AI tool took an action outside its approved scope — sent an email, submitted a form, accessed a file, or modified a record without authorization.

Fabricated output acted upon: AI-generated output containing false or fabricated information was used to inform a Town decision, communication, or action before the error was caught.

Vendor breach: A vendor whose AI platform processes Town data suffers a security breach, exposing Town information.

Unauthorized recording: An AI meeting assistant recorded a conversation without the consent of all parties, potentially violating M.G.L. ch. 272, § 99.

Bias or discrimination: AI output used in a resident-facing decision is found to reflect bias against a protected class.

9.2 Response Steps

Step	Action	Responsible Party
1. Contain	Stop the AI tool immediately. Revoke any active access, OAuth connections, or API keys. Do not delete evidence.	Person who discovered the incident
2. Document	Record: what happened, which AI tool was involved, what data was affected, who may be impacted, and when the incident was discovered. Document the incident using the format described in this section.	Person who discovered the incident + their manager
3. Notify	Notify the Town Administrator and the IT Provider within 24 hours. If the incident involves potential exposure of personal information of Massachusetts residents, proceed to Step 3a.	Manager / Department Head
3a. Breach Notification	If personal information (as defined in M.G.L. ch. 93H) of Massachusetts residents may have been compromised, the Town must: (1) notify the Massachusetts Attorney General's office, (2) notify the Director of Consumer Affairs and Business Regulation, and (3) notify affected individuals. Consult Town Counsel immediately.	Town Administrator + Town Counsel
4. Investigate	Determine root cause, scope of impact, and whether the incident is ongoing. Engage the IT Provider for technical investigation. Preserve all logs and records.	IT Provider + Town Administrator
5. Remediate	Implement corrective actions: policy changes, additional training, tool removal, access changes, vendor	Town Administrator

	engagement, or disciplinary action as appropriate. Document all remediation steps.	
6. Review	Conduct a post-incident review within 30 days. Update this policy if the incident reveals a gap. Report findings to the Select Board if the incident involved resident data or public impact.	Town Administrator

⚠ MASSACHUSETTS DATA BREACH NOTIFICATION — M.G.L. ch. 93H

Massachusetts law requires any entity that owns, licenses, stores, or maintains personal information of Massachusetts residents to provide notice of a breach of security. 'Personal information' includes a resident's first and last name (or first initial and last name) in combination with a Social Security number, driver's license number, financial account number, or other specified data elements. If an AI incident results in the unauthorized acquisition or use of such information, the Town has a legal obligation to notify the AG's office and affected residents as soon as practicable and without unreasonable delay. Consult Town Counsel immediately in any situation where this may apply.

SECTION 10: Training Requirements

Reading and acknowledging this policy is necessary but not sufficient. Training ensures that all covered individuals understand not just the rules, but the reasoning behind them and how to apply them in practice.

10.1 Required Training

Training Type	Who	When	Content
Initial Training	All covered individuals	Within 30 days of policy adoption or onboarding	Full policy overview, approved tools, data handling rules, pre-use checklist, traceability requirements, incident reporting
Annual Refresher	All covered individuals	Annually, on or before the anniversary of policy adoption	Policy updates, new tools, lessons learned from incidents, emerging AI risks, refresher on data handling and ethical output
Manager / Team Lead Training	Managers and team leads	Within 30 days of assuming a supervisory role	'Requires Review' approval procedures, traceability oversight, how to evaluate AI use cases, vendor risk review basics

Agentic Tool Training	Anyone who uses or approves agentic tools	Before first use of any agentic tool	Agentic tool risks, persistent access management, action log review, prompt injection awareness, incident response for autonomous actions
Elected Official / Committee Member Briefing	Select Board members, appointed committee members	Within 30 days of taking office or appointment	Policy overview with emphasis on Open Meeting Law, public records, and executive session restrictions

10.2 Training Records

The Town Administrator (or designee) maintains a training log recording: the name of each covered individual, the date training was completed, and the type of training.

Training completion is a condition of continued access to sanctioned AI tools. Covered individuals who have not completed required training may have their access suspended until training is completed.

Training records are retained per the Massachusetts Municipal Records Retention Schedule.

SECTION 11: Roles, Enforcement & Review

11.1 Roles & Responsibilities

Role	Responsibilities
Town Administrator (Policy Owner, in coordination with IT Provider)	Maintains and updates this policy. Approves new AI tools (with IT Provider input). Receives all escalations. Conducts or commissions annual review. First point of contact for gray-area questions. Maintains training log.
All Covered Individuals (Employees, Elected Officials, Committee Members, Volunteers, Contractors)	Read and acknowledge this policy annually. Follow all rules in this policy. Use the Pre-Use Checklist (Section 2.2). Report Shadow AI tools and potential violations. Complete required training within specified timelines.
Managers / Team Leads	Enforce this policy within their teams. Approve 'Requires Review' use cases. Ensure traceability records are kept. Authorize and periodically review any agentic tool access within their team.
Elected Officials & Committee Members	Comply with this policy when using AI in connection with Town business. Ensure AI use in public meetings complies with the Open Meeting Law (Section 5.2). Follow executive session rules in Section 5.2. Disclose AI involvement in motions, reports, or

	recommendations presented to the public.
IT Provider	Maintains the Approved Tools list (Section 12). Monitors for unsanctioned AI usage on Town devices and networks. Reviews and documents all persistent access grants for agentic tools per Section 6.2. Alerts Town Administrator when new AI features appear in existing tools. Supports vendor risk assessments and incident response.

11.2 Enforcement

Severity	Example	Response
Low — Unintentional	Covered individual uses an unapproved browser AI extension for internal tasks; no Sensitive Data involved.	Coaching conversation + tool review
Medium — Negligent	An approved meeting tool was left running inadvertently after a meeting ended or before consent was confirmed; resident data entered into a personal AI account; AI output sent externally without review.	Formal reminder + required training + incident log. If resident data was exposed, initiate AI Incident Response (Section 9).
High — Deliberate or Damaging	Knowingly activating a recording tool after a participant declined consent, or using an unsanctioned tool to record without disclosure; knowingly entering regulated data into an unsanctioned tool; agentic tool granted access beyond approved scope; AI output sent externally with fabricated facts.	Formal review up to termination. Possible resident notification and legal review. If personal information was compromised, mandatory breach notification under M.G.L. ch. 93H. If a meeting was recorded without all-party consent, potential criminal liability under M.G.L. ch. 272, § 99 (up to 5 years imprisonment, \$10,000 fine). Town Counsel consulted in all High-severity cases.

11.3 Policy Review

This policy is reviewed annually at a minimum.

An unscheduled review is triggered when: a new AI tool materially changes how covered individuals work; a significant AI-related incident occurs; relevant Massachusetts or federal regulations change; or a major platform update alters how a sanctioned tool handles data.

Persistent access reviews for agentic tools are conducted per Section 6.2.

Version history is maintained by the Town Administrator. Previous versions will be provided upon request.

11.4 Procurement Integration

Any AI tool purchase, subscription, or free-tier adoption must follow the Town's existing procurement policy in addition to the Tool Approval Process in Section 4.3. The AI policy review does not replace procurement — both are required. Specifically:

If the tool requires a purchase or subscription, follow the Town's standard procurement process (including any dollar-threshold approvals) before submitting the AI tool review request.

If the tool is free but processes Town data, it still requires AI policy review under Section 4.3.

The IT Provider should be consulted during procurement to assess technical compatibility, security requirements, and integration risks.

Contract language for AI tool vendors should include data handling, training-on-inputs, breach notification, and termination/data-deletion provisions.

11.5 Vendor Risk Assessment Checklist

The following checklist is used by the Town Administrator and IT Provider when evaluating an AI tool for approval under Section 4.3. All items should be documented before a tool is approved.

#	Assessment Criteria	Status
1	Data Residency: Where is data processed and stored? Is it within the United States?	☐
2	Training on Inputs: Does the vendor use customer inputs (prompts, data) to train or improve its AI models? Can this be disabled?	☐
3	Breach Notification: What are the vendor's contractual obligations to notify the Town in the event of a data breach? Does the timeline meet Massachusetts requirements?	☐
4	Data Retention & Deletion: How long does the vendor retain customer data? Can the Town request deletion? What happens to data upon contract termination?	☐
5	Subprocessors: Does the vendor share data with subprocessors or third parties? Are they disclosed?	☐
6	Security Attestation: Does the vendor hold SOC 2 Type II, ISO 27001, or equivalent security certification?	☐
7	Insurance: Does the vendor maintain cyber liability insurance coverage?	☐
8	Terms of Service: Do the Terms of Service grant the vendor any rights to customer data beyond what is necessary for service delivery?	☐
9	Access Controls: Does the tool support role-based access, SSO, and audit logging?	☐
10	Regulatory Compliance: Does the vendor comply with applicable regulations (e.g., FERPA, HIPAA, CJIS) relevant to the Town's use case?	☐

SECTION 12: Approved AI Tools

The following tools have been reviewed and approved for use by covered individuals of the Town of Chester. Each tool is approved only for the uses and under the conditions described below. Using an approved tool outside its stated conditions requires a new review under Section 4.3.

Tool	Account Type	Approved Use	Conditions / Restrictions
ChatGPT	Team account (organizational)	Internal drafting, brainstorming, summarization, research	No Sensitive Data. No resident, employee, or financial data. Do not use personal accounts. All output must be reviewed before use.
Microsoft Copilot (Outlook)	Organizational license (M365)	Email drafting, summarization, reply suggestions	Review all AI-generated responses before sending. Do not use for resident-facing communications without a second review. Data stays within M365 tenant.
Google Gemini (Chrome)	Organizational account	Research, web summarization, general productivity	No Sensitive Data input. Do not use for resident-facing content without department head review. Verify all facts before use.
Microsoft Teams AI Notetaker	Organizational license (M365)	Meeting transcription and summarization	ALL participants must consent before activation per M.G.L. ch. 272, § 99. Do not use in executive sessions without Town Counsel approval. Output stored on authorized Town systems only. Transcripts handled as Sensitive Data.

MAINTAINING THIS LIST

This list is maintained jointly by the Town Administrator and IT Provider. Tools are added through the approval process in Section 4.3. Removed tools are documented with the date and reason for removal. All covered individuals are notified when the list changes.

SECTION 13: Acknowledgment

By signing below, I confirm that I have read and understood this AI Governance Policy (Version 2.1) and agree to comply with its requirements.

Field	
Name	
Role / Title	
Department / Board / Committee	
Signature	
Date	
Training Completed	☐ Yes ☐ No
Training Date	

— End of Policy —