GREATER SALT LAKE
Municipal Services District

Greater Salt Lake Municipal Services District Technology Policy Handbook

Document Purpose

This handbook provides a comprehensive framework for managing, securing, and governing Greater Salt Lake Municipal Services District (MSD) technology systems and services. It is designed to ensure legal compliance, safeguard public resources, and promote effective use of information technology within municipal agencies. Each policy is aligned with applicable Utah state laws and best practices in public sector technology management.

Policy Owner: IT Department
Effective Date: [Insert Date]
Last Reviewed: [Insert Date]
Approved By: [Insert Authority]

Acceptable AI Usage Policy

1. Purpose

Defines acceptable and responsible use of AI technologies in compliance with Utah law and federal standards.

2. Scope

Applies to all employees, contractors, and third-party use of AI tools and systems within the organization.

3. Legal and Regulatory Framework

- Utah Senate Bill 149 (2024)
- Utah Code § 63A-16
- Utah Code § 13-44
- Utah Code § 63G-2
- NIST AI RMF

4. Definitions

The following definitions are provided to ensure all staff have a clear understanding of terms used in this policy. Both formal definitions and plain-language explanations are included.

A. Artificial Intelligence (AI)

Formal Definition:

Artificial Intelligence (AI) refers to the field of computer science and technology focused on creating systems that can perform tasks that typically require human intelligence. These tasks may include decision-making, problem-solving, learning, pattern recognition, language understanding, and prediction. AI encompasses a broad range of technologies, from rule-based automation to advanced machine learning algorithms.

Plain-Language Explanation:

AI is when we use computers or software to do things that normally need human thinking, like recognizing speech, recommending products, or detecting fraud. Think of it as teaching a computer to 'think' in certain ways to make our jobs easier.

B. Generative Artificial Intelligence (Generative AI)

Formal Definition:

Generative Artificial Intelligence (Generative AI) is a subset of AI that uses machine learning models to create new content such as text, images, audio, video, or code. Unlike traditional AI, which classifies or predicts based on existing data, generative AI produces novel outputs that resemble human-created work. Examples include text-generation systems, image synthesis tools, and audio/video generation technologies.

Plain-Language Explanation:

Generative AI can actually create things, like writing an article, drawing a picture, or making music. It doesn't just look at data; it makes something new based on what it has learned. Tools like ChatGPT or image generators are common examples.

C. Large Language Model (LLM)

Formal Definition:

A Large Language Model (LLM) is a type of generative AI trained on vast amounts of text data to understand and generate human language. LLMs use advanced deep learning techniques to predict and produce contextually relevant text responses, answer questions, summarize content, translate languages, and perform other natural language processing tasks. Common examples include models such as GPT (Generative Pre-trained Transformer) and similar transformer-based architectures.

Plain-Language Explanation:

An LLM is a special kind of AI that's really good at working with words. It can write emails, summarize reports, answer questions, or even translate text. Think of it as a very advanced 'smart assistant' that understands and generates language.

5. Acceptable Uses of AI

This section defines acceptable and permitted uses of Artificial Intelligence (AI) technologies within the organization. These examples are intended to guide staff in understanding how AI can be appropriately applied to support daily work while ensuring compliance with organizational values, ethical standards, and legal requirements.

A. Automation

AI tools may be used to automate repetitive, time-consuming, or low-risk tasks. Examples include processing data entries, generating routine system reports, sorting emails, and monitoring system performance. Automation helps improve efficiency and allows staff to focus on higher-value responsibilities.

B. Scheduling

AI systems may be used to assist with scheduling meetings, reminders, and resource allocation. For example, AI-enabled calendar assistants can find optimal meeting times across teams or departments, improving collaboration and reducing time spent on coordination.

C. Education and Training

AI tools may be used to support employee education, training, and professional development. This includes personalized learning platforms, tutoring systems, training simulations, and AI-generated learning resources. These tools must be reviewed to ensure accuracy, fairness, and alignment with organizational learning goals.

D. Internal Summaries

AI may be used to generate internal summaries of documents, reports, and meeting notes to support productivity. These summaries are intended for internal use only and should not replace official records or externally shared communications unless reviewed and approved by staff.

E. Communications

MSD Communications staff may use approved artificial intelligence (AI) tools to support official duties, including drafting, editing, summarizing, or organizing written materials; developing ideas for communication campaigns or social media content; and improving clarity, readability, and consistency in public messaging. All AI-generated content must be reviewed and verified for accuracy, tone, and compliance with MSD standards, policies, and procedures prior to use or publication. AI tools are intended to enhance efficiency and creativity, not to replace human expertise or professional judgment. Staff must ensure that all use of AI aligns with MSD's values of transparency, accuracy, and public trust, and that no confidential, sensitive, or proprietary information is entered into AI platforms.

E. Translation

AI-powered translation tools may be used to translate internal documents, messages, or resources to improve accessibility and communication among staff and stakeholders. Translations used for official or public-facing purposes must be reviewed by a qualified human translator to ensure accuracy.

F. Chatbots and Virtual Assistants

AI-powered chatbots and virtual assistants may be used to provide internal support, answer frequently asked questions, and assist with routine processes such as password resets, policy inquiries, or IT support requests. These tools must be configured to protect user privacy and should not be used to handle sensitive, regulated, or confidential information without appropriate safeguards.

All permitted uses of AI must follow organizational policies for data security, privacy, compliance, and oversight. Staff must exercise judgment when using AI tools and seek guidance if unsure about an AI application's appropriateness.

6. Prohibited or Restricted Uses

This section defines prohibited or restricted uses of Artificial Intelligence (AI) technologies within the organization. These restrictions are necessary to ensure AI use is ethical, secure,

legally compliant, and aligned with organizational values. Staff must not use AI in the following ways:

A. Handling Sensitive or Regulated Data

AI tools must not be used to process, analyze, or store sensitive, confidential, or regulated data without explicit authorization and safeguards. Examples include personally identifiable information (PII), financial data, medical records, and information protected under laws such as HIPAA, FERPA, or GDPR.

B. Decision-Making Without Human Oversight

AI must not be used as the sole decision-maker in areas that significantly affect employees, stakeholders, or the public. This includes hiring, disciplinary actions, financial approvals, legal interpretations, or safety-critical decisions. Human review and accountability are required. This review includes but is not limited to the following: state and federal statute compliance, engineering plan building code compliance, citizen problem responses, and document search result accuracy and completeness.

C. External Communications Without Review

MSD Communications staff are prohibited from using AI tools to create, distribute, or approve content that is false, misleading, confidential, or inconsistent with MSD policies or public information standards. AI tools may not be used to impersonate individuals, produce unauthorized statements, or make determinations that require management or legal review. Staff must not input personal, private, or restricted information into AI systems that are not approved for secure data handling. The use of AI must comply with all applicable laws, copyright protections, data privacy requirements, and ethical communication practices. AI-generated text, images, or translations must not be used for official, public-facing, or legal communications without prior human review and approval. Examples include press releases, legal documents, policy statements, and public announcements.

D. Misinformation or Manipulation

AI must not be used to intentionally generate or spread misinformation, disinformation, deepfakes, or other content designed to mislead, manipulate, or cause harm. This includes impersonating individuals, creating fraudulent content, or misrepresenting organizational positions.

E. Discrimination or Bias

AI systems must not be used in ways that reinforce, amplify, or introduce unfair bias or discrimination. Staff must be vigilant when using AI outputs in contexts that may affect equity, fairness, or compliance with anti-discrimination laws.

F. Unauthorized Third-Party Tools

Staff must not use unauthorized or unvetted third-party AI tools, particularly those that involve uploading internal data or documents to external platforms. Only AI solutions approved by the organization's IT and leadership teams may be used.

Violations of these prohibited or restricted uses may result in disciplinary action, up to and including termination, and could carry legal consequences. When in doubt, staff should seek approval from the IT department before using AI in sensitive or uncertain contexts..

7. Transparency and Disclosure

AI-generated content must be labeled; public-facing tools must disclose AI interaction.

8. Data Protection and Model Input Controls

PII, PHI, CJIS, or confidential data may not be input into public models.

9. Procurement and Development Requirements

AI tools must pass a risk review, with contracts ensuring transparency, control, and data protection. See Appendix A for specific review criteria.

10. Documentation and Record Retention

Logs of AI use and datasets must be retained per GRAMA standards.

11. Training and Awareness

All users must complete AI usage training before accessing approved AI systems.

12. Roles and Responsibilities

IT: Enable secure tools
IT Director: Approve use cases
Legal: Ensure compliance
Users: Follow policy

13. Enforcement and Incident Handling

Violations may result in access revocation, HR discipline, or referral to state authorities.

14. Policy Review

Annual review or after regulatory changes or AI-related incidents.

Appendix A: Risk Review Criteria Examples

1. Privacy and Data Protection Risk View

- **Description**: Evaluates whether AI tools collect, process, or expose personally identifiable information (PII), protected health information (PHI), or other sensitive government records.
- **Context**: For example, an AI chatbot handling resident inquiries may inadvertently store Social Security numbers, driver's license information, or addresses without adequate safeguards.
- **Risk Mitigation**: Require data minimization, encryption, strict access controls, and compliance with Utah Government Records Access and Management Act (GRAMA) and federal privacy laws (e.g., HIPAA if applicable).

---

2. Bias, Fairness, and Equity Risk View

- **Description**: Reviews whether AI outputs unintentionally discriminate or produce biased results against protected classes (race, gender, age, disability, etc.).
- **Context**: In the hiring processes, AI screening tools could disproportionately filter out minority or older candidates; or in public safety, predictive tools may unfairly target certain neighborhoods.
- **Risk Mitigation**: Require human oversight, bias testing, and policy alignment with Title VI of the Civil Rights Act and ADA accessibility requirements.

---

3. Operational and Accountability Risk View

- **Description**: Assesses risks related to reliability, explainability, and accountability of AI-driven decisions.
- **Context**: If a city uses AI to classify zoning applications or prioritize service requests (e.g., pothole repairs), errors or unexplained rejections can create legal liability and erode public trust.
- **Risk Mitigation**: Require human-in-the-loop review, audit trails, clear documentation of AI decision logic, and assignment of accountability to staff; not the AI system.

AI Data Privacy Policy

1. Purpose

Establishes privacy requirements for AI systems to ensure compliance with Utah data privacy laws and responsible data use.

2. Scope

Applies to all employees, vendors, and systems handling data processed by AI across the organization.

3. Legal and Regulatory Framework

- Utah Consumer Privacy Act (UCPA)
- Utah Code § 13-44 (Personal Information Protection)
- GRAMA (Utah Code § 63G-2)
- Utah SB149 (AI Policy Act)

4. Principles of AI Data Privacy

The following principles must be followed whenever Artificial Intelligence (AI) systems are designed, implemented, or used within the organization. These principles ensure compliance with legal obligations, ethical standards, and the protection of individual rights.

A. Data Minimization

AI systems should only collect and process the minimum amount of personal or sensitive data necessary to achieve the stated purpose.

Key Point: Avoid unnecessary data collection or retention.

Example: If an AI tool is used to schedule meetings, it should not request unrelated information such as home addresses.

B. Purpose Limitation

Data used by AI must only be applied for the specific, legitimate purpose for which it was originally collected. Any secondary use requires approval and, where necessary, renewed consent. Approval is required in writing from the General Manager, Head of IT and the department Head with their consultation with legal counsel in complex cases. Consent must be requested in writing from the parties the information is collected from.

Key Point: Do not reuse data for unrelated tasks without authorization.

Example: Data collected for training an internal chatbot should not later be used for marketing purposes without explicit approval.

C. Consent and Transparency

Individuals must be informed when their data is being collected, how it will be used in AI systems, and whether AI tools are involved in decision-making. Where legally required, explicit consent must be obtained.

Key Point: Users and employees should understand how AI is using their data.

Example: A notice should inform employees if AI is monitoring workflow efficiency or analyzing emails for productivity insights.

D. Data Accuracy

Data used by AI must be accurate, up to date, and relevant to the task. Processes should be in place to correct errors and prevent the use of outdated or misleading information.

Key Point: Inaccurate data can lead to faulty or unfair AI outcomes.

Example: If an AI system is used to screen job applications, incorrect applicant information must be corrected promptly to avoid biased decisions.

E. Fairness and Non-Discrimination

AI must be designed and used in ways that avoid unfair bias or discrimination against individuals or groups. Regular reviews should be conducted to identify and mitigate potential bias in training data, algorithms, or outputs.

Key Point: AI decisions should be equitable, inclusive, and respectful of diversity.

Example: An AI-based loan approval system must not disadvantage applicants based on gender, race, or other protected characteristics.

These principles must be applied across all AI projects to protect individual rights, ensure ethical use, and maintain trust in organizational AI systems.

5. Restrictions on Use of Personal and Regulated Data

- No PII/PHI/CJIS entered into public or multi-tenant LLMs
- De-identify or synthesize data where possible

6. Third-Party AI Vendors and Cloud Services

- Must comply with Utah Code § 13-44-202
- Require breach notice, data control, audit rights

7. Model Training and Dataset Governance

- Maintain dataset inventory and documentation
- Review datasets for bias, consent, and compliance

8. AI System Design and Privacy by Design

- Implement PETs, access controls, and audit logs

9. Individual Rights and DSARs

- Support access, correction, deletion requests per UCPA

10. Audit and Compliance

- Conduct PIAs, log access, and audit AI data use

11. Roles and Responsibilities

IT: Enable secure tools
IT Director: Approve use cases
Legal: Ensure compliance
Users: Follow policy

12. Enforcement

Violations may result in discipline, termination, or regulatory notification

13. Policy Review

Reviewed annually or after regulatory/tech changes or incidents

Human Oversight of AI Policy

1. Purpose

Ensures all AI systems are governed by appropriate levels of human oversight, accountability, and intervention to protect rights and trust.

2. Scope

Applies to all AI systems, employees, vendors, and departments interacting with or deploying AI solutions.

3. Legal and Ethical Framework

- Utah SB149 (2024)
- Utah Code § 13-44, § 13-61
- GRAMA – Utah Code § 63G-2
- NIST AI RMF

4. Key Definitions

Human oversight of AI systems refers to the processes, practices, and safeguards put in place to ensure that artificial intelligence (AI) technologies operate in ways that are ethical, accountable, safe, and aligned with organizational and societal values. Oversight ensures that AI systems support, rather than replace, human judgment in areas where outcomes may impact rights, safety, fairness, or compliance with laws and regulations.

5. Oversight Principles

A. Accountability: Humans remain ultimately responsible for AI-driven decisions and outcomes, particularly in high-risk or sensitive contexts.

B. Decision Review: AI output must be reviewed, verified, or approved by qualified individuals before being used in critical or legally binding decisions.

C. Intervention and Control: Humans must retain the ability to intervene, override, or stop an AI system if it produces harmful, biased, or incorrect outputs.

D. Transparency: AI systems should provide explainable results that allow human reviewers to understand how outputs were generated.

E. Risk Management: Oversight includes continuous monitoring of AI systems for errors, misuse, bias, or unintended consequences, with corrective actions applied as needed.

F. Ethical Safeguards: Oversight ensures AI systems do not compromise fairness, non-discrimination, safety, or compliance with applicable laws and organizational policies.

Practical Examples of Human Oversight

- In recruitment, AI may screen job applications, but a human hiring manager must review final candidates to ensure fairness and compliance with employment laws.

- In engineering, AI may assist in analyzing design data, but professionals must make the final recommendation.

- In finance, AI may identify unusual transactions, but compliance officers must validate whether flagged transactions are fraudulent before action is taken.

In summary, human oversight of AI systems ensures that AI technologies remain tools that enhance, rather than replace, human judgment. It establishes a safeguard against risks, maintains accountability, and protects individual rights while enabling the benefits of AI innovation.

6. Human Oversight Requirements by Risk Tier

Low – Document and basic audit trails
Moderate – Periodic human review
High – Mandatory HITL, justification logging, appeal rights

7. Oversight Mechanisms

- Pre-deployment review
- Ongoing monitoring and intervention
- Fallback systems and documentation

8. Oversight in Automated Decision-Making

- Designated human decision-maker
- Right to appeal
- No fully automated decisions without approval

9. Vendor and External AI Systems

- Must demonstrate oversight capabilities
- Provide logs and allow audits

10. Oversight Roles and Responsibilities

Governance: Define risk tiers
Departments: Enforce oversight

IT/Data: Enable rollback and logging
Legal: Ensure compliance

11. Training and Competency

- AI ethics training
- Scenario-based oversight simulations
- Familiarity with AI model limitations

12. Audit and Reporting

- Log human decisions and interventions
- Review logs during annual AI audits

13. Enforcement and Violations

- Model suspension, investigation, breach reporting per Utah Code § 13-44-202 (See Appendix A)

14. Policy Review

Annual or after high-risk AI deployment or incident

Appendix A - Summary of Utah Code § 13-44-202

Key Obligations under Utah Code § 13-44-202

1. Investigation Upon Breach (akin to Model Suspension for AI)

When a person or entity "owns or licenses" computerized data containing personal information of Utah residents, and becomes aware of a **breach of system security**, they must:

- Conduct a **reasonable and prompt investigation**, in good faith, to determine **whether personal information has been or likely will be misused for identity theft or fraud**.

This aligns with the concept of **suspending** the AI model or system by halting further use while evaluating the incident.

2. Notification to Affected Utah Residents

If the investigation reveals that misuse **has occurred or is reasonably likely to occur**, the entity must:

- **Notify each affected Utah resident**, using one of these methods:

- First-class mail to the most recent address;
- Electronically, if that's the primary communication method and consistent with federal rules (e.g., under 15 U.S.C. § 7001);
- Telephone (including lawful auto-dialing);
- If none of the above are feasible: publish notice in a **newspaper of general circulation** and pursuant to statutory requirements for public notices.

3. Escalated Reporting to Government & Agencies

Depending on the number of affected Utah residents, additional reporting is required:

- **If 500 or more residents** are affected (or likely to be), the entity must also notify:
  - The **Utah Office of the Attorney General**
  - The **Utah Cyber Center**
- **If 1,000 or more residents** are affected (or likely to be), the entity must further notify:
  - **Each nationwide consumer reporting agency** (e.g., credit bureaus)

4. Timing and Coordination Requirements

Notifications must be issued in the "**most expedient time possible without unreasonable delay**," taking into account these factors:

- Legitimate investigative needs of law enforcement;
- Scope assessment of the breach;
- Restoration of reasonable system integrity.

Additionally:

- If a **law enforcement agency** requests a delay in notifying affected parties (to avoid impeding a criminal investigation), the entity may delay, but must send notifications **immediately once law enforcement confirms it's safe to proceed**.

5. Third-Party Cooperation

If an entity **maintains data but does not own or license it**, upon discovery of a breach (with misuse having occurred or likely to occur), they must:

- **Inform and cooperate** immediately with the **owner or licensee** of the data; sharing relevant breach information.

6. Protected Status of Government Reports

Reports submitted to the Attorney General or Utah Cyber Center under the 500+ threshold may be treated as **confidential protected records**, provided they meet statutory conditions. These protected submissions may include:

- The date the breach occurred;
- The date it was discovered;
- The number of people affected (including Utah residents);
- The type of personal information involved;
- A brief description of the incident.

7. No Waiver Permitted

Any attempt to contractually waive compliance with this section is **unenforceable and void**.

Summary Table

Stage / Action	Requirement
Investigation (Suspension)	Promptly investigate breach in good faith for misuse risk.
Resident Notification	Notify affected Utah residents if misuse occurred or is likely.
Government Reporting	Notify AG & Utah Cyber Center if ≥ 500 affected; also consumer reporting agencies if ≥ 1,000.
Timing	Notify as soon as practicable; may delay only at law enforcement request and then resume promptly.
Third-Party Coordination	Data custodians must notify and cooperate with owners/licensees if breach discovered.
Report Confidentiality	Submissions to AG/Cyber Center may be protected; include key breach details.
Waiver Prohibition	Any waiver of these requirements is void as against public policy.

Practical Takeaway for AI Governance in a City Government

If your city operates AI systems that process Utah residents' personal information, you should ensure your **acceptable use policy and incident response protocols** include:

1. **Immediate suspension or containment** of the AI system upon a suspected breach.
2. **Prompt, documented investigation** of misuse risk.
3. **Tiered notification procedures**:
    - Residents,
    - Utah AG & Cyber Center (≥ 500 affected),
    - Consumer reporting agencies (≥ 1,000 affected).
4. **Clear protocols** for delayed notification only when law enforcement mandates.
5. **Obligations for third-party processors** to notify city IT/data owners.
6. **Secure handling of breach reports** including confidentiality for higher-impact incidents.
7. **Ensuring no policy or contract allows waiving these obligations.**

AI Oversight Checklist

This checklist supports the Human Oversight of AI Policy and ensures accountability, fairness, and legal compliance for AI deployments within the organization.

1. Governance and Approval

- ☐ Has the AI system been classified by risk tier (Low, Moderate, High)?
- ☐ Has the AI Governance Committee approved deployment (if Moderate or High)?
- ☐ Is the AI system listed in the official AI Model Inventory?

2. Human Oversight Design

- ☐ Is there a designated human reviewer for AI-assisted decisions?
- ☐ Are intervention and override mechanisms in place?
- ☐ Has a fallback plan been documented in case the AI system fails or malfunctions?

3. Transparency and Disclosure

- ☐ Are users or affected individuals informed when AI is used?
- ☐ Are the AI model's decisions explainable to internal stakeholders?
- ☐ Is the AI vendor or internal team providing documentation on model behavior?

4. Training and Competency

- ☐ Have oversight personnel completed AI ethics and governance training?
- ☐ Are staff trained on how to intervene and override AI decisions?

5. Logging and Auditing

- ☐ Are logs being captured for all interventions and AI decisions?
- ☐ Are logs reviewed regularly by IT Security or Audit?
- ☐ Are audit results shared with the AI Governance Committee?

6. Legal and Policy Compliance

- ☐ Does the AI use comply with Utah SB149, UCPA, and GRAMA?
- ☐ Is there a process for individuals to appeal or challenge AI decisions?
- ☐ Have contracts with AI vendors included human oversight and audit clauses?

Third-Party AI Integration Policy

1. Purpose

Defines standards for integrating third-party AI systems while ensuring compliance with Utah law, transparency, and accountability.

2. Scope

Applies to all departments, contractors, and vendors integrating third-party AI into the organization. AI capabilities include, but are not limited to: natural language generation or understanding; text, image, audio, or video generation; speech-to-text and text-to-speech; machine translation; summarization; classification, clustering, recommendation, or ranking; anomaly or fraud detection; predictive analytics and forecasting; computer vision (detection, recognition, segmentation); autonomous agents or copilots; reinforcement learning systems; and any automated decisioning or decision support using machine learning or statistical learning methods

3. Legal and Regulatory Framework

- Utah SB149 (2024)
- Utah Code § 13-61, § 13-44, § 63G-2 (GRAMA)
- NIST AI RMF

4. Third-Party AI Definition

"Externally developed AI tool, API, model, or cloud-based service offering AI capabilities" means any software application, software-as-a-service (SaaS) platform, infrastructure service, application programming interface (API), software library, model artifact (including foundation, base, fine-tuned, or hosted models), or managed service that delivers or exposes artificial intelligence functionality and is created, owned, hosted, or maintained by an entity other than the Organization. This includes commercial, open-source, free, freemium, trial, beta, community, or research offerings regardless of cost, hosting location, or deployment model.

Ownership and Hosting Considerations

The term applies whether the capability is accessed via public cloud, vendor-managed private cloud, on-premises appliance delivered by a vendor, edge device managed by a vendor, or downloaded/open-source artifacts (models, weights, libraries) that are obtained from third parties and used internally.

Inclusions (Examples)

• Hosted foundation models or APIs (e.g., general-purpose LLM endpoints, image-generation APIs).

• Vendor-provided copilots, chatbots, or assistants embedded in productivity, CRM, ERP, HRIS, or ITSM suites.

• Open-source model weights, checkpoints, or inference servers obtained from public registries or research labs, even if self-hosted by the Organization.

• Third-party SDKs, libraries, and plug-ins that enable AI features inside Organization-built applications (e.g., vector search SDKs, embedding libraries, on-device ML kits).

• Managed AutoML, feature stores, model hosting, prompt orchestration, RAG services, and agent frameworks provided by external vendors.

• Analytics or marketing platforms that incorporate AI-powered targeting, scoring, or personalization.

• Security tools that use ML/AI for detection, classification, or response (e.g., anomaly detection, behavior analytics).

Exclusions (Non-Examples)

The term does NOT include:

• AI systems wholly designed, trained, and maintained internally by the Organization using only Organization-created code and data, with no third-party model artifacts, endpoints, or libraries involved.

• Traditional deterministic automation or rule-based scripts without any ML/AI component (e.g., simple macros, if/then task scripts).

• Non-AI cloud services that do not expose or rely upon AI/ML features.

Note: If any third-party AI library, model artifact, or hosted inference component is incorporated, the system becomes an externally developed AI capability for policy purposes.

Edge Cases and Clarifications

• Fine-Tuned Models: If an external base model is fine-tuned by the Organization, the resulting model remains in scope as externally developed due to dependence on the third-party base.

• Open-Source: Open-source models or libraries sourced from outside the Organization are considered external regardless of internal hosting.

• Embedded AI in Products: If a purchased product includes AI features (even if optional), those features are covered by this definition.

• Shadow AI: Trial, personal, or free-tier use of external AI tools by staff for work purposes is covered and must follow procurement, security, and privacy review requirements.

• Data Residency/Transfer: External status applies even if the vendor claims in-region hosting; data handling and cross-border transfer risks remain subject to review.

• Offline or On-Device Models: Externally sourced models running locally (e.g., on laptops or mobile devices) are in scope.

Policy Hooks (Applicability of Controls)

Use of externally developed AI capabilities triggers the Organization's requirements for:

• Third-party risk assessment, security and privacy review, and data protection impact assessment where applicable.

• Contractual safeguards (e.g., data use restrictions, IP ownership, service levels, incident response, model update/rollback controls).

• Compliance checks for regulated data types, records retention, accessibility, and audit logging.

• Human oversight, accuracy testing, bias and fairness evaluation, and explainability commensurate with risk.

• Inventory and registration in the AI system catalog before production use.

Quick Reference Examples

In Scope: LLM API from a cloud provider; open-source vision model downloaded and fine-tuned; vendor chatbot built into a help desk platform; AutoML service for forecasting.

Out of Scope: In-house rule-based script; Organization-built model trained from scratch with only Organization code/data and no external components.

5. Pre-Integration Risk Assessment

- Evaluate purpose, data, bias, transparency, security, and compliance

6. Contractual Requirements for AI Vendors

- Scope of use and human accountability
- Data use limitations and security standards
- Bias and audit clauses
- Breach notification and audit rights

7. Prohibited Integrations

- Black-box AI without transparency
- Unencrypted sensitive data usage
- Full automation without oversight
- UCPA or SB149 violations

8. Ongoing Oversight and Monitoring

- Monitor for drift, security issues, or contract violations
- Conduct annual compliance review

9. Transparency and Disclosure

- Notify public users of AI use
- Identify provider and role
- Provide point of contact for appeals

11. Roles and Responsibilities

Procurement: Contract terms
IT: Technical checks
Legal: Compliance review
Departments: Outcome monitoring

12. Policy Violations

- Suspension, termination, breach notification
- Legal or disciplinary action

13. Policy Review

- Annual review
- After major integration or incident
- Upon legal or regulatory updates

Use of AI in Video and Audio Meetings Policy

1. Purpose

To regulate the use of AI technologies in video and audio meetings to ensure compliance with Utah laws and protect sensitive data.

2. Scope

Applies to all employees, contractors, and vendors involved in virtual meetings with or for the organization.

3. Legal and Regulatory Framework

- Utah SB149 (2024)
- Utah Code § 13-61 – UCPA
- Utah Code § 13-44 – Protection of Personal Information
- Utah Code § 63G-2 – GRAMA
- NIST AI RMF

4. Permitted AI Tools

Only built-in AI tools from approved platforms such as Microsoft Teams, Google Meet, and Zoom Workplace are permitted for use during video and audio meetings. These tools are considered secure, compliant, and supported by the organization. Any use of non-approved third-party AI meeting tools is prohibited unless the IT department grants an explicit exception on a per-meeting basis.

Rationale for Restricting Other Tools

While third-party AI transcription, summarization, or analytics tools may appear convenient, they present several risks to the organization and its stakeholders. The following points explain why limiting AI tool usage to approved, built-in platforms is necessary:

A. Data Privacy and Security Risks: Many external AI tools transmit meeting content—including confidential discussions, personally identifiable information, or sensitive business data through external servers where the organization has no control over storage, retention, or access.

B. Regulatory and Legal Compliance: Use of unapproved AI tools may violate state, federal, or contractual requirements for data handling. This is especially critical where laws such as HIPAA, FERPA, or GDPR apply.

C. Unauthorized Data Retention: Third-party tools may store transcripts, recordings, or summaries indefinitely, creating risks of data leakage or unintended disclosure.

D. Accuracy and Reliability: Built-in platform tools undergo rigorous testing and vendor support, whereas third-party applications may produce inaccurate or misleading transcripts or summaries, which could affect decision-making.

E. Vendor Risk and Accountability: External AI tools may not provide clear accountability, support, or guarantees. If sensitive information is leaked or misused, the organization may face reputational damage and legal liability.

F. Integration and Support Issues: Built-in AI features are maintained and updated by approved vendors, ensuring compatibility with enterprise systems and IT support. Unsupported tools can create conflicts or technical failures during meetings.

By restricting AI usage in video and audio meetings to only built-in tools from approved platforms, the organization ensures data security, compliance with legal requirements, and consistent reliability of services. Exceptions will only be considered when justified, reviewed, and formally approved by the IT department.

5. Prohibited AI Tools and Services

- Otter.ai
- Fireflies.ai
- Trint
- Descript
- Sonix.ai
- Unapproved browser extensions

These tools are banned due to privacy, control, and compliance concerns.

6. Responsibilities

IT: Configure and monitor tools
Legal: Ensure legal compliance
Meeting Hosts: Notify participants
Employees/Vendors: Use only permitted platforms

7. Consent and Disclosure Requirements

Participants must be notified at meeting start.
Stored transcripts must comply with data retention and public record laws.

8. Exceptions

Allowed only with:
1. Legal/business justification
2. Approval from IT and Legal
3. Proof of compliance

9. Monitoring and Enforcement

- IT monitors for violations
- Disciplinary action for internal misuse
- Vendor contract termination if applicable
- Breach reporting per Utah Code § 13-44-202

10. Policy Review and Updates

Reviewed annually, after regulatory changes, or following an incident.

Regulatory Framework Summaries

Utah Guidance

- **SB 149 (2024) – AI Policy Act**
  Requires disclosure when generative AI is used in public-facing contexts; sets limits for regulated professions. Cities must ensure staff and vendors disclose AI use.
- **Utah Code § 63A-16 – Technology Governance**
  Establishes state CIO/DTS standards (security, accessibility, procurement). Cities should mirror governance practices for consistency with state systems.
- **Utah Code § 13-44 – Data Breach Notification**
  Requires prompt investigation of breaches, resident notification, and escalated reporting to AG, Utah Cyber Center, and credit agencies at scale. AI systems handling PII must comply.
- **Utah Consumer Privacy Act (UCPA) – § 13-61**
  Grants Utah residents rights to access, delete, and opt-out of data uses. While cities are often exempt, **vendors serving cities are subject**. Cities should require vendor compliance.
- **GRAMA – § 63G-2 (Government Records)**
  Governs classification, retention, and disclosure of records. AI prompts, logs, and outputs may qualify as public records.

Federal / National Guidance

- **NIST AI Risk Management Framework (RMF 1.0)**
  Voluntary but widely recognized. Provides a four-part cycle (Govern, Map, Measure, Manage) to ensure AI is reliable, safe, and explainable.

Implications for City Leaders

- **Policy Alignment** – Adopt disclosure, privacy, and records rules aligned with state law.
- **Vendor Management** – Require contractors to comply with UCPA, SB 149, and breach-notification duties.
- **Risk Oversight** – Incorporate NIST RMF and EO principles into procurement and governance.

- **Public Trust** – Demonstrate transparency, fairness, and accountability in AI use.