02.7 Artificial Intelligence Security

This policy section establishes Eversource's security, risk, and governance requirements for the use of Artificial Intelligence (AI), including 1) approved Eversource AI Platform tools such as Microsoft 365 Copilot and Microsoft 365 Copilot Chat 2) in-house generative Artificial Intelligent (AI) systems, such as Regulatory Knowledge Engine (RKE), and 3) custom agentic solutions. The intent is to enable responsible innovation while protecting Company data, ensuring regulatory compliance, and reinforcing human accountability in all AI assisted outcomes.
This policy builds upon the existing Risk Management framework and aligns with NIST 800-171 while reflecting Eversource approved AI guardrails and governance processes.

02.7.01 Artificial Intelligence Handling

Establishes requirements for artificial intelligence (AI).
This applies to:

- Use of AI tools (local, embedded, web based, or API driven)
- Custom agents developed on an approved Eversource platform, such as Microsoft 365 Copilot, Copilot Studio, or similar platforms
- Third-party or external AI services
- In-house generative artificial intelligence (AI) systems, such as RKE or RA
- Connectors, plugins, APIs, or integrations that enable AI tools to access Company data

This applies to AI use in all Eversource environments, which includes but is not limited to all, local development, sandbox, non-production, and production environments.

Control Standard
Audit and Compliance of Artificial Intelligence (AI)
EECS-1017
Regular audits must be conducted to verify compliance with this policy and NIST 800-171 requirements.

Control Standard
Data Access Controls for Artificial Intelligence (AI)
EECS-1018
Access to confidential or restricted information used in generative AI systems, must be restricted to authorized personnel only.

Control Standard
Data Classification of Artificial Intelligence (AI)
EECS-1019
All data used in generative artificial intelligence (AI) systems must be classified according to its level of sensitivity. Data classification terms and handling requirements referenced in this policy (e.g., Public, Internal Use, Confidential, Highly Restricted, and regulated data categories) must align to and be interpreted consistently with Eversource's corporate data classification standards. Where definitions differ, the corporate data classification standard is authoritative.

The following data must not be entered, uploaded, or processed in any AI system unless explicitly approved by the Enterprise AI Team:

- Highly Restricted information
- Confidential Business Information (CBI)
- Protected or Confidential Personal Information (PPI/CPI)
- Critical Infrastructure Protection (CIP) data
- Critical Energy Infrastructure Information (CEII)
- Security credentials, keys, or authentication data
- Attorney Client Privileged Information

Approved and Restricted AI Tools

Approved Eversource Platform Gen AI tools:

- Microsoft 365 Copilot
- Microsoft 365 Copilot Chat
- Microsoft Foundry – advanced models

Public or third-party AI tools without Enterprise level agreements and unique Eversource tenants are considered public AI systems. Access to these tools and use of company data in these tools is prohibited unless explicitly reviewed and approved by Information Security and the Enterprise AI Team.

Acceptable Use (Allowed vs. Prohibited)

Allowed:

- Use approved Eversource AI tools for business productivity (e.g., drafting, summarizing, brainstorming, and analysis) using data that is appropriate for the tool and permitted by Company data classification requirements.
- Use AI to improve or accelerate work outputs only when a qualified human validates accuracy, completeness, and appropriateness before use or distribution.

Prohibited:

- Entering, uploading, or processing Highly Restricted, Confidential Business Information (CBI), Protected/Confidential Personal Information (PPI/CPI), CIP data, CEII, or security credentials/keys in any AI system unless explicitly approved by the Enterprise AI Team.
- Using public AI systems, unapproved third-party AI tools, or non-enterprise AI tenants for any Company data or Company work.
- Sharing AI outputs as authoritative when the output has not been independently validated (including calculations, policy interpretations, or regulatory guidance) or checked for hallucinations.
- Using AI to execute automated actions, initiate workflows, or make operational decisions without required Enterprise AI Team approval and Human-in-the-Loop oversight.
- When interfacing with AI agents, agents should not automatically inherit all end user permissions.

Tooling Controls (Installation and Use of AI Tools)

Installation or use of any AI tool, browser extension, plug-in, connector, API integration, or client application that is not explicitly approved and provided through an Eversource-licensed enterprise offering is prohibited. Employees must use only the AI tools and capabilities provisioned through approved Eversource platforms, and licensed tenants made available by the Enterprise AI Team and/or Information Technology. This applies to personal, team, and enterprise agents as well. All agents must be built and used only within approved Eversource platforms.

General Agentic AI Use

Agentic AI Approved for Use:

- Microsoft 365 Copilot Agents built by Microsoft, such as Researcher and Analyst
- Custom read-only personal or team agents built using Microsoft 365 Agent Builder (see more below)
- In-house generative artificial intelligence (AI) agents built and monitored by the Enterprise AI Team

Eversource allows the use and custom development of agents. Risk and approval requirements increase as agent autonomy, reach, and connectivity to systems increase.
Personal Agent (individual use):

- Built for an individual's personal productivity within the Eversource ecosystem using approved Eversource platforms and licensed tenants
- Must use personal data only and must not access, ingest, store, or process Company data unless explicitly approved by the Enterprise AI Team.
- Must not be shared enterprise wide or with external parties.

Team Agent (workgroup use):

- Built for a defined team or workgroup within the Eversource ecosystem using approved Eversource platforms and licensed tenants
- Shared only within the defined team/workgroup and must follow least-privilege access and existing Eversource access controls.
- May access Company data only through approved Eversource platforms and only to the extent permitted by existing access controls and data classification requirements.

Enterprise Agent (broad organizational use):

- Built for broad distribution (multiple teams, departments, or enterprise-wide) within the Eversource ecosystem using approved Eversource platforms GenAI
- Requires Enterprise AI Team review and approval prior to deployment, sharing, or enabling access.
- Must meet enterprise governance requirements for security, privacy, monitoring, and change management, and must enforce least-privilege access and data classification controls.

Agentic AI Requiring Explicit Approval
Custom built agents require review and approval by the Enterprise AI Team if they:

- Complete tasks or take automated actions
- Integrate with third-party systems (i.e. Workday) beyond Microsoft applications
- Access Confidential or Highly Restricted data (see section EECS-1019 for more details)
- Are shared beyond an immediate team

Control Standard

Engagement with the Enterprise AI Team

EESC-1030
AI use-cases that involve connectors, automation, agents, or access to company data must not be pursued independently.
Employees must engage the Enterprise AI Team early when exploring or implementing AI use cases that introduce additional risk, complexity, or system integration. These activities must not be undertaken independently.
Engagement with the Enterprise AI Team is required when a use case includes:

- Connectors, plugins, APIs, or integrations
- Automation or workflow execution

- Agents shared beyond an immediate team
- Use of Confidential or regulated data
- Any uncertainty about whether a use case is permitted

This engagement ensures appropriate security, governance, and risk oversight before development or deployment.

Control Standard
Generative Artificial Intelligence (AI) Models
EESC-1024
Microsoft and OpenAI are our approved model providers. Within OpenAI, only their GPT models are approved for general use. Use of any additional or alternative models requires prior review and approval from the Enterprise AI Team.

Control Standard
Human-in-the-Loop (HITL) Oversight for AI Systems
EECS-1025
AI systems must include a qualified human reviewer to validate, override, or audit AI generated outputs before final action is taken. This includes, but is not limited to, developer code and content created using AI systems.

Control Standard
In-House Artificial Intelligence (AI) Model Training
EECS-1020
In-House Artificial Intelligence (AI) Model Training EECS-1020 Generative Artificial Intelligence (AI) models developed and maintained in-house may be trained using confidential or restricted information, with approval from the data owner(s) and the office of the Chief Information Security Officer (CISO). In-house AI models (AI systems developed, implemented, and maintained within the company, where input and generated data used for training and operation remains within the company's control and does not leave the organization) must be designed and implemented to ensure that all data remains within the company's control. Technical controls must be employed to protect the confidentiality, integrity, and availability of the data, including but not limited to:

- Data Encryption: All confidential or restricted information used for training in-house AI models must be encrypted both at rest and in transit. Encryption algorithms and key management procedures must adhere to industry best practices and comply with applicable regulatory requirements and Eversource policies.
- Access Control: Strict access control mechanisms must be in place to limit access to confidential or restricted information used for training in-house AI models. Access should be granted based on the principle of least privilege and role-based access control (RBAC), ensuring that only authorized personnel with a legitimate need can access the data.
- Data Segregation: Confidential or restricted information used for training in-house AI models should be segregated from other data sources and stored in secure, dedicated environments to minimize the risk of unauthorized access or data leakage.

- Secure Development Practices: In-house AI model development must follow secure software development practices, including regular security testing, vulnerability scanning, and code review to minimize potential security risks and ensure compliance with NIST 800-171, OWASP, and Eversource security requirements.
- Logging and Monitoring: All activities related to the use of confidential or restricted information for training in-house AI models must be logged and monitored for suspicious activities or potential policy violations. The logs must be securely stored and periodically reviewed by the appropriate security personnel.

Control Standard
Model Output Monitoring for Artificial Intelligence (AI)

EECS-1028

Outputs generated by in-house generative artificial intelligence (AI) systems must be monitored and reviewed to ensure that they do not inadvertently disclose confidential or restricted information. In the event of such disclosure, appropriate actions must be taken to address the issue, including retraining the model or adjusting the system's parameters.

AI usage may be logged, monitored, and audited. Non-compliance may result in revocation of access, remediation requirements, or disciplinary action under HR policy. Use of unapproved or unauthorized AI tools (including installation of non-approved AI software, extensions, plug-ins, or integrations) may result in immediate revocation of access and required remediation. Repeated or intentional non-compliance may be escalated under applicable HR and Information Security processes.

Exceptions to this policy must be formally approved through the established GRC exception process.

Control Standard
Third-Party Generative Artificial Intelligence (AI) Systems

EECS-1023

Third-party AI tools, connectors, plugins, APIs, or integrations are not permitted unless formally reviewed and approved via established governance processes and by the Enterprise AI Team. Any AI solution that connects to external systems or data sources must undergo existing security, privacy, and risk review through established enterprise governance processes for approval.

Control Standard
Training and Awareness for Artificial Intelligence (AI)

EECS-1029

Employees, contractors, and stakeholders involved in the development, implementation, or use of generative artificial intelligence (AI) technology must receive appropriate training on Eversource AI policies and NIST 800171 requirements. Training must be conducted periodically to ensure that all relevant personnel understand their responsibilities related to public AI systems, third-party AI tools, and in-house AI models. Engagement with the Enterprise AI Team supports governance and risk review and does not replace required AI training and awareness activities

Control Standard
Validation of AI Output Accuracy
EECS-1026
AI systems must be integrated with accuracy-check mechanisms, such as performance metrics, confidence scores, or corroboration with known data sources.
AI generated outputs must not be assumed to be accurate. Users are responsible for independently validating factual accuracy, calculations, and policy or regulatory interpretations.
AI generated confidence levels, probabilities, or inferred insights must not be treated as authoritative and may not be used for formal decisions without corroboration from trusted sources.

Revision History

July, 2025
- AI Policy updates

July, 2024
- Policy Update for cloud compliance.

January, 2024
- Policy Update for compliance with NIST 800-17

December, 2023
- Policy Template Update for compliance with NIST 800-17

July, 2023
- Created new section for Artificial Intelligence Security

May, 2022
- Updated data classification to include Highly Restricted
- Added applicable NIST 800-171 standards

May, 2021
- Original publication of policy, standards and applicable control procedures.