Maryland
DEPARTMENT OF INFORMATION TECHNOLOGY

Wes Moore | Governor
Aruna Miller | Lt. Governor
Katie Savage | Secretary

MEMORANDUM FOR: All State of Maryland Executive Branch Staff
FROM: Katie Savage, Secretary of the Department of Information Technology and Chair of the Governor's AI Subcabinet
SUBJECT: Responsible Artificial Intelligence (AI) Policy
VERSION: 1.0
EFFECTIVE DATE: 05/23/25
CHANGELOG: N/A

Maryland
DEPARTMENT OF INFORMATION TECHNOLOGY

Wes Moore | Governor
Aruna Miller | Lt. Governor
Katie Savage | Secretary

The State of Maryland's Responsible AI Policy

Purpose

This policy provides a governance framework for AI systems used by or on behalf of the State of Maryland (the “State”), enabling State agencies to use AI systems for the benefit of constituents while safeguarding against potential harms.

The key objectives of the AI Policy are to:

- Provide guidance that is clear, easy to follow, and supports decision-making for those who may be purchasing, configuring, developing, operating, or maintaining the State's AI systems or leveraging AI systems to provide services to the State's constituents.
- Ensure adherence to the State's guiding principles with regards to how AI systems are purchased, configured, developed, operated, or maintained.
- Define roles and responsibilities related to the State's usage of AI systems.
- Establish and maintain processes to assess and manage risks presented by AI systems used or considered for use by the State.
- Align the governance of AI systems with existing data governance, security, and privacy measures in accordance with the State's Information and Security Policy and the State's Data Policy.
- Define prohibited uses of AI systems.
- Establish "sunset" procedures to safely retire AI systems that no longer meet the State's needs or requirements.
- Define how AI systems may be used for legitimate State purposes in accordance with applicable state and federal laws, and existing agency policies.

Scope

This policy applies to:

1. All AI systems and use cases deployed (or under consideration) by State executive branch agencies; and
2. Staff (full-time, part-time, contractual), interns, consultants, contractors, partners, and volunteers who may be purchasing, configuring, developing, operating, or maintaining the State of Maryland's AI systems or who may be leveraging AI systems to provide services to the State.

Although this Responsible AI Policy establishes the State's AI policies, the AI Implementation Guidance outlines how agencies should operate with regard to AI systems in alignment with the State's AI policies. The State's AI systems and the data contained therein will be purchased, configured, developed, operated, and maintained using the State's AI Implementation Guidance.

Terms & Definitions

Artificial Intelligence: "Artificial intelligence" or "AI" is a machine-based system that can, for a given set of human-defined objectives, make predictions, recommendations, or decisions influencing real or virtual environments. Artificial intelligence systems use machine- and human-based inputs to perceive real and virtual environments; abstract such perceptions into models through analysis in an automated manner; and use model inference to formulate options for information or action.

Algorithm: A series of logical steps through which a computer or software program turns particular inputs into particular outputs.

AI system: Any system, software, sensor, or process that automatically generates outputs including, but not limited to, predictions, recommendations, or decisions that augment or replace human decision-making. This extends to software, hardware, algorithms, and data generated by these systems, used to automate large-scale processes or analyze large data sets.

Proof of Concept (PoC): A test, evaluation, demonstration, or pilot project of a technology in a controlled and scoped environment to see if it can be successfully and responsibly deployed to benefit the State.

Agentic Artificial Intelligence (Agentic AI): A system designed with the capability to proactively and autonomously pursue defined goals within a specific environment (real or virtual).

Harm: NIST's AI Risk Management Framework (AI RMF) defines AI harm as encompassing impacts on individuals, organizations, and ecosystems, including harm to civil liberties, rights, physical/psychological safety, economic opportunities, reputation, business operations, and global systems.

Guiding Principles for Responsible AI Systems

These principles describe the State's values with regards to how AI systems are purchased, configured, developed, operated, or maintained. These principles are drawn and expanded from Governor Moore's AI Executive Order.

1. **Human-Centered Design:** AI systems are developed and deployed with a human-centered approach that evaluates AI powered services for their impact on the public and innovation potential for State services and resident outcomes.

2. **Security & Safety:** AI systems maintain confidentiality, integrity, and availability through safeguards that prevent unauthorized access and use. AI systems must be reliable and safe, and minimize risks to individuals, society, and the environment.
3. **Privacy:** Privacy is preserved in all AI systems by safeguarding Personal Information (PI) and sensitive data from unauthorized access, disclosure, and manipulation. Please reference the State Data Classification Policy for further guidance.
4. **Transparency:** The purpose and use of AI systems is proactively communicated and disclosed to the public. An AI system, its data sources, operational model, and policies that govern its use are understandable and documented.
5. **Equity:** AI systems support equitable outcomes for everyone. Bias in AI systems is proactively identified and mitigated to ensure equitable outcomes and minimize harm to any individuals impacted by their use.
6. **Accountability:** Clear roles and responsibilities govern the deployment and maintenance of AI systems, and human oversight ensures adherence to relevant laws and regulations.
7. **Effectiveness:** AI systems are reliable, meet their objectives, and deliver precise and dependable outcomes for the contexts in which they are deployed.

Policy

When purchasing, configuring, developing, operating, or maintaining AI systems, the agency will:

1. uphold these Guiding Principles for Responsible AI Systems, State and federal laws, orders, and regulations;
2. obtain technical documentation (e.g., the specific model used by the AI system, cyber risk mitigations taken by the vendor) about AI systems as part of due diligence and oversight;
3. surface and apply relevant risks, mitigations, and guardrails;
4. require contractors to comply with the requirements of the State's Responsible AI Policy;
5. add AI use cases to the AI inventory once in production or use; and
6. in the event of an incident involving the use of the AI system, an agency will follow its Incident Response Plan and report the incident to DoIT.

Roles & Responsibilities

- The **Department of Information Technology (DoIT)** is responsible for:
    - Maintaining and continually improving the intake process and supporting policies and guidance in consultation with relevant cybersecurity, data, privacy, and legal teams.
    - Providing technical assistance and policy guidance through the intake process.

    - Consultation on “buy vs build” decisions and evaluation of AI technology for responsible use, with relevant stakeholders as needed.
    - Providing guidance on the management of "high-risk" AI systems.
    - Annually aggregating and publishing the State AI Inventory based on agency submissions.
    - Generating “Governance Cards” which provide insight into how to responsibly use AI for particular use cases, in addition to recommendations for AI tools to power these use cases.
- The Governor's **Artificial Intelligence Subcabinet**:
    - Oversees the implementation of Maryland's AI strategy and governance.
    - Coordinates decisions with state-wide ramifications.
    - Provides informed recommendations to the Governor on AI-related matters.
- **Executive Agencies** are responsible for:
    - Appointing an AI lead for their agency who is responsible for ensuring that the relevant DoIT and agency policies, guidance, and best practices in shepherding agency use cases from concept to production are followed. These AI leads will coordinate and work alongside their agency’s Portfolio Officers, Data Officers, and Privacy Officers.
    - Taking through the AI Intake process any new or modified AI system use case.
    - Implementing risk mitigation strategies, impact assessments, and disclosures surfaced during intake and through procurement.
    - Retiring or modifying AI systems that fail to meet ongoing requirements or have been designated for "sunset."
    - Implementing approaches to continuous monitoring or auditing of high-risk AI systems.
    - Submitting “live” AI agency use cases to the AI Inventory.

AI Risk Classification System

Maryland follows a risk-based approval approach for AI projects. Projects are classified into into one the following risk tiers by agencies during intake, with different tiers corresponding to different levels of oversight and due diligence:

1. **Unacceptable Risk:** Systems posing extreme risks to public welfare, safety, or rights that cannot be mitigated. These include AI and/or Agentic AI uses that violate fundamental rights (e.g., unlawful surveillance or unchecked social scoring of citizens). Such systems are banned entirely from use.
2. **High-Risk:** AI systems that may significantly affect individuals or critical government operations. These could be systems that influence decisions on health, safety, law

enforcement, eligibility for essential services, privacy, financial or legal rights, or other high-impact outcomes. High-risk AI is permissible only with robust safeguards, including a comprehensive AI Risk Assessment before deployment, implementation of risk mitigation measures, and ongoing monitoring.

Data Level Alignment: Systems will be considered high-risk if they meet the data classification criteria of Level 3 (Confidential) or 4 (Restricted) Data. Refer to the Responsible AI Implementation Guidance document for approaches on applying these classifications.

3. **Limited Risk:** AI systems that have moderate or low impact and few inherent risks. These may include AI tools that improve internal efficiency or customer service without making autonomous decisions that affect constituents. Such systems are allowed provided they have been submitted through the DoIT intake process.

   Data Level Alignment: These systems have moderate/low impact, often improving internal efficiency without making autonomous decisions. If they meet the data classification criteria of Level 1 (Public) or 2 (Protected/Internal Use Only) Data, refer to the Responsible AI Implementation Guidance document for approaches on applying these classifications.

4. **Minimal Risk:** AI applications that pose a negligible risk and are tools for internal use or infrastructure with no impact on individual safety or rights. These require no special approval beyond standard DoIT intake process.

   Data Level Alignment: These applications pose a negligible risk and don't impact individual safety or rights. This aligns best with the following data classification criteria of Level 1 (Public) or 2 (Protected/Internal Use Only) Data refer to the Responsible AI Implementation Guidance document for approaches on applying these classifications.

Prohibited Uses

The use of certain AI systems is prohibited due to the sensitive nature of the information processed and severe potential risk. This includes the following prohibited purposes:

- Real-time and covert biometric identification: The live identification of an individual using technologies including, but not limited to, facial recognition and iris scanning without that individual's knowledge or meaningful consent.

- Emotion analysis: The use of computer vision techniques to classify human facial expressions, body movements, or language into emotions or sentiments (e.g., positive, negative, neutral, happy, angry).
- Fully automated decisions (e.g., Agentic AI) that specifically fall under the use cases of "Unacceptable Risk" outlined above.
- Social scoring: The use of AI systems to track and classify individuals based on their behaviors, socioeconomic status, or personal characteristics.
- Cognitive behavioral manipulation of people or specific vulnerable groups.

If State staff become aware of an instance where an AI system has caused harm, staff must report the instance to their supervisor, AI agency lead, and DoIT (via intake).

Sunset Procedures

If an AI system operated by the State or on its behalf ceases to provide a positive utility (e.g., found to have become an "Unacceptable Risk" application due to a change in usage or intent of the tool) to the State's residents, then the use of that AI system must be halted unless an express exception is provided by the head of the agency, following a thorough review of potential impacts on services and stakeholders and the AI Subcabinet informed of this decision by the head of the agency. If the abrupt cessation of the use of that AI system would significantly disrupt the delivery of State services, the agency may maintain the services but must present a remediation plan to DoIT and the agency's leadership. Additionally, any impacted individuals must be provided with notification of the finding and the ability to opt out of such a system.

Public Records and Transparency

The State is committed to public transparency in its use of AI. Data stored and processed in AI Systems is subject to the Maryland Public Information Act. DoIT will aggregate an AI Inventory annually, based on each agency's submissions, and make it publicly available. Use cases related to ensuring the safety and security of State systems will be excluded if disclosing this information could compromise system security or integrity. State staff must follow all current procedures for records retention and disclosure.