VIRGINIA IT AGENCY

COMMONWEALTH OF VIRGINIA

Enterprise Architecture (EA)

Information Technology Resource Management (ITRM)

Artificial Intelligence (AI) Utilization Policy

EA Policy (EA400)
June 2026

Preface

Publication Designation
Enterprise Architecture Policy (EA400)

Subject
Enterprise Architecture

Effective Date
(Date)

Supersedes
Past versions (see below table)

Scheduled VITA Review
Periodically or as needed

Authority

Code of Virginia, §2.2-2006
(Definitions)

Code of Virginia, §2.2-2007
(Powers of the CIO)

Code of Virginia, §2.2-2007.1
(Additional duties of the CIO relating to information technology planning and budgeting)

Code of Virginia, § 2.2-2009
(Additional duties of the CIO relating to security of government information)

Code of Virginia, § 2.2-2012
(Additional powers and duties related to the procurement of information technology)

Code of Virginia, §2.2-603(F)
(Authority of agency directors, with respect to IT and data security and risk management)

Scope
This standard is applicable to all Executive Branch state agencies and institutions of higher education (hereinafter collectively referred to as "agencies") that are responsible for the management, development, purchase, and use of information technology resources in the Commonwealth of Virginia. This standard does not apply to research projects, research initiatives or instructional programs at public institutions of higher education.

Purpose
This standard establishes *the framework for enterprise architecture* direction and technical requirements, which govern the acquisition, use and management of information technology resources by executive branch agencies.

General Responsibilities

Chief Information Officer of the Commonwealth (CIO)
Agency head of VITA. Responsible for and approves statewide technical and data policies, standards, guidelines, and requirements for information technology, including with respect to information technology planning, procurement, and security.

Virginia Information Technologies Agency (VITA)
At the direction of the CIO, VITA leads efforts that draft, review, and update technical and data policies, standards, guidelines, and requirements for information technology.

VITA uses requirements in IT technical and data related documents when establishing contracts; reviewing procurement project, security and budget requests and strategic plans, and when developing and managing IT enterprise and infrastructure services.

Executive Branch Agencies
Provide input and review during the formulation, adoption and update of statewide technical and data policies, standards and guidelines for information technology.

Comply with the requirements established by COV policies and standards. Apply for exceptions to requirements when necessary.

Related ITRM Policies, Standards, and Guidelines
Enterprise Architecture Standard (EA225)

Reviews

Updates to this publication and opportunities for review occur through the regulatory process for guidance documents.

Publication Version Control

Please direct questions related to this publication to VITA's Enterprise Architecture Division (EA) at ea@vita.virginia.gov. VITA notifies the Agency Information Technology Resources (AITRs) at all state agencies, institutions, and other interested parties of revisions to this document.

The following table contains a history of the revisions to this publication.

Version	Date	Revision Description
1.0	5/26/2024	Original
1.1	6/30/2026	Additional guidance on disclaimers. Added Agency Governance section. Textual revisions and reformatting for readability and clarity.

Table of Contents

Preface **2**
Reviews 3
Publication Version Control 3

Table of Contents **4**

Overview **5**

Policy Standards **5**
Ethical Use of AI 5
Business Case for AI 5
AI Approval Process 5
AI Disclaimers 6
Mitigating Third Party Risks 7
Protecting Citizens' Data 7
Agency Governance 8

Overview

The Commonwealth of Virginia (COV) is well-positioned to lead in the artificial intelligence (AI) sector. The Commonwealth's strengths — a robust data center industry, strong higher education institutions, a growing workforce, and a state government seeking greater efficiency — make it a natural innovator in this space.

AI is already in use across many Commonwealth agencies for data processing, automated decision-making, customer service, and operational efficiency. As that use continues to grow, the public deserves a clear understanding of how AI is employed on their behalf. It is therefore essential to establish comprehensive, consistent standards for the responsible, ethical, and transparent use of AI by all Commonwealth agencies.

Policy Standards

Ethical Use of AI

These guiding principles shall be used to ensure that AI is developed and used responsibly.

- Every Department, Agency, and Office is responsible for ensuring that its use of AI — including generative AI — is trusted, safe, secure, ethical, and transparent.
- AI models must be fully documented and made available for review.
- AI generated outcomes must be validated by humans for bias and unintended consequences.
- All Commonwealth entities must ensure their AI use is resilient, accountable, and explainable. "Black box" AI — systems whose decision logic cannot be explained — shall not be used in any decision-making or approval process.

Business Case for AI

A Commonwealth Department, Agency, or Office may deploy generative AI or other AI capabilities only when there is a clear, positive outcome for Virginia's citizens. Qualifying purposes include:

- Reducing wait times for government services
- Removing barriers to access for government institutions and services
- Reducing bureaucratic delays
- Lowering the cost — in time and money — of interacting with government for individuals and businesses
- Improving the quality and delivery of government services
- Making Virginia a safer and more productive Commonwealth for all residents, regardless of location or socioeconomic status

AI must be the optimal solution for the stated purpose. Before selecting an AI approach, agencies must examine alternative technology or process solutions. This examination must include a Regulatory Impact Analysis (RIA) that addresses costs and benefits, data sources and methods, digital sovereignty, and alternatives to the proposed change.

A clear statement of the AI application's intended purpose — including whether AI will make a recommendation to a user or an autonomous decision on a user's behalf — must be included in the COV AI registry.

AI Approval Process

To ensure AI is used in a trusted, safe, and secure manner, all Commonwealth Departments, Agencies, and Offices must complete the following approval process before deploying any AI system — whether for internal or external use.

Definitions

- Internal AI systems: Solutions using generative AI or related capabilities exclusively within a

department, agency, or office to improve internal efficiency or operations. They do not include any system that produces a decision affecting an individual citizen or business.
- External AI systems: Solutions that analyze data about individual citizens or businesses, make decisions affecting those individuals or businesses, or produce output directly accessible by citizens or businesses.

Approval Workflow
Any internal or external AI system that an agency seeks to develop, implement, or procure must be entered into the COV AI Registry and complete the following sequential approval process with the following reviewers:

- Agency IT Representative (AITR) and Information Security Officer (ISO)
- VITA and the CIO of the Commonwealth
- Secretary of Administration

VITA shall retain records of all AI uses and approvals and shall ensure that timely notifications are sent to the agency and agency Secretariat upon submission and upon each approval decision.

Approval Criteria
At a minimum, the manager responsible — and, where applicable, the AITR and ISO — must evaluate the following before approving any AI system:

- Fairness: Verify the system will not result in unlawful discrimination or disparate impact on any individual or group based on age, genetic information, color, ethnicity, race, creed, religion, national origin, ancestry, sex, gender identity or expression, sexual orientation, marital status, familial status, pregnancy, veteran status, disability, or lawful source of income.
- Public benefit: Confirm that the AI capability will benefit citizens and advance the agency's objectives within the intended context of use.
- Human oversight: Assess the degree to which humans are involved in monitoring and validating AI outputs.
- Risk assessment: Identify inherent risks — including cybersecurity, data privacy, and risks to individual or business health and safety — and confirm steps have been taken to mitigate them. Document any additional guardrails being put in place.
- Data stewardship: Confirm appropriate stewardship of Commonwealth-held data.
- Cost-benefit analysis: Complete a cost impact analysis covering data sources, methods, and alternatives to the proposed approach.
- Vendor warranties: Assess whether the AI developer provides — or should provide — warranties or assurances regarding safety, security, cybersecurity resilience, and output reliability. Document any gaps and steps taken to address them.

Exemptions
This approval process does not apply to:

- AI used to defend Commonwealth security systems
- AI embedded within standard commercial software products
- AI research and development activities or instructional programs at public institutions of higher education

AI Disclaimers
The Commonwealth requires public disclosure whenever generative AI or other AI capabilities are used in any process, decision, or output that affects citizens or businesses. Disclosure is required in the following situations:

- Consumer engagement: Individuals must be informed when they are interacting with an AI system rather than a human.

- Pre-use: Individuals must be notified before AI is used to make a decision about them.
- Content labeling: AI-generated images, video, audio, or text must be clearly labeled as such.
- Human resources: Applicants and employees must be informed when AI is used in hiring or employment decisions.
- Public website: Visitors must be informed when AI is used to deliver services on a public-facing website.

Required Disclaimer Language

- When AI generates a decision or output directly:
  DISCLAIMER: This decision or output was generated by artificial intelligence.
- When AI assists in a broader process but does not make the final decision:
  DISCLAIMER: This decision or output was created with assistance from artificial intelligence.
- When AI is used to produce content:
  DISCLAIMER: This (content) was generated with the assistance of artificial intelligence.

Transparency Requirements for External AI Systems

AI systems making external decisions affecting Commonwealth citizens must also:

- Disclose how AI is used to arrive at a decision
- Describe the extent of human involvement in validating and overseeing those decisions
- Clearly list any available options for individuals to appeal those decisions

Mitigating Third Party Risks

To address cybersecurity, data privacy, and misuse risks associated with third-party AI vendors, all Commonwealth Departments, Agencies, and Offices must rigorously vet any third-party AI developers, system administrators, providers, or contractors. Required actions include:

- Vendor selection: Evaluate each vendor's ability to deliver value to the Commonwealth and its citizens, and assess their trust, safety, and security practices against industry best practices.
- Data standards compliance: Ensure vendors apply industry-standard best practices for data collection and use, including protection of personally identifiable information (PII) and compliance with all applicable Commonwealth laws and regulations.
- Testing review: Review results of any testing or red-teaming conducted by the vendor — including efficacy, cybersecurity, physical security, risk, and bias testing — and work with vendors to address identified issues before implementation.
- Documentation review: Review audit reports, product roadmaps, warranties, terms of service, end-user license agreements, contracts, and other vendor documentation to support value assessment and risk management.
- Inventory management: Maintain an inventory of all third-party materials (hardware, open-source and proprietary software, foundation models, open-source and proprietary data) used in AI systems.
- Security screening: Verify that third-party AI resources and personnel undergo security audits and screenings. Failure to provide relevant security information is a risk indicator.
- Watermarking: Apply watermarking technologies to Commonwealth materials produced by generative AI as a deterrent to data and model extraction attacks.
- Legal and ethical compliance: Ensure all AI systems align with applicable laws, regulations, and guidelines, and conduct regular reviews to maintain compliance as technology and best practices evolve.
- Staff training: Educate government employees and decision-makers about the benefits and risks of AI, including awareness of potential biases and operational challenges.

Protecting Citizens' Data

The Commonwealth prioritizes privacy and data protection as agencies develop, implement, and procure AI systems. All AI systems must comply with the Government Data Collection and Dissemination Practices Act (§

2.2-3800 et seq.). In addition, all Commonwealth offices and agencies must:

- Data minimization: Use only the data necessary for the AI system's stated purpose. AI systems must not have unrestricted access to large volumes of personal data.
- Data security and retention: Secure all data and retain it only as long as necessary to fulfill the AI system's objective. Establish data retention timeframes where feasible.
- Anomaly monitoring: Monitor for anomalies using approaches such as control limits, confidence intervals, and integrity constraints.
- Security metrics: Establish and track AI system security tests and metrics, including red-teaming activities, frequency of anomalous events, system downtime, incident response times, and time-to-bypass.
- User controls: Implement proper user controls so individuals know when their data is being used to produce AI outputs or automated decisions.
- User consent: Allow users to consent to the use of their data by AI systems whenever possible.
- Sensitive data handling: Sensitive, confidential, and protected data may only be used in private AI systems exclusively accessible to Commonwealth users.
- AI data governance accountability: Each agency using AI systems shall designate a responsible official for AI data governance compliance with both this policy and the GDCDPA. Agencies shall conduct annual audits of high-risk AI systems to verify compliance with data minimization requirements, proper consent documentation, maintenance of access records, and fulfillment of data subject requests. Audit results shall be submitted to VITA and the Secretary of Administration.
- Enhanced protections for vulnerable populations: AI systems processing data concerning children, individuals receiving public benefits, health and medical information, or criminal justice require enhanced protections including additional human oversight and validation, heightened security measures beyond standard safeguards, and regular audits for bias and disparate impact. Biometric data (fingerprints, voiceprints, faceprints, retinal scans, iris patterns) used in AI systems must employ encryption at rest and in transit.
- De-identification and privacy protection: Agencies shall apply NIST-compliant de-identification standards to AI training data and utilize differential privacy techniques for aggregated data analysis where appropriate. Re-identification attempts on de-identified data are prohibited. When de-identified data is used in AI systems, agencies shall maintain documentation of the de-identification methodology and assess re-identification risks
- Third-party AI vendor requirements: All contracts with third-party AI vendors must prohibit vendors from using Commonwealth citizen personal information for training or improving commercial AI products, require compliance with all GDCDPA requirements, mandate data deletion within 90 days of contract termination unless retention is required by law, and prohibit data sharing with vendor affiliates or partners without explicit approval.

Agency Governance

Agencies are accountable for overseeing AI use both within their departments and by the individuals who access them. Agencies shall:

- Maintain an AI Use Form documenting that employees, contractors, and authorized users have read and acknowledged this policy
- Document where and how AI is used across their IT portfolio
- Ensure that their COV AI Registry entries are complete and current
- Monitor their use of AI solutions, tracking usage volume and patterns, user adoption, quality and effectiveness, technical performance, cost, and alignment with the originating business case

All governance materials shall be available to VITA as requested.