



Governance and Oversight of Artificial Intelligence in Virginia

Study resolution

- Review use, governance, regulation and oversight of AI in Virginia*
 - AI use in state government, including opportunities and risks
 - whether agency staff have appropriate skills to use AI effectively
 - whether state AI policies and processes enable innovation while mitigating risks
 - approaches to help the General Assembly identify and address the most substantial risks to the public from AI, beyond state government use
 - key factors that affect the cost and availability of AI systems

*JLARC study resolution, authorized December 2025. Higher ed use of AI will be addressed in 2027 JLARC study pursuant to HJ 32 (2026).

Primary research activities

- Interviewed key stakeholders in Virginia and nationally
 - VITA leadership and staff
 - IT leadership at state agencies of varying size and mission
 - DLS, JCOTS, OAG, and legislative staff in other states
 - Industry representatives (e.g., Google, Microsoft, Amazon, Anthropic)
 - National experts
- Surveyed senior IT leaders at all executive branch and independent agencies (100% response rate)

VITA = Virginia Information Technologies Agency · DLS = Division of Legislative Services · JCOTS = Joint Commission on Technology and Science · OAG = Office of the Attorney General

Primary research activities (cont.)

- Analyzed state employee AI usage data and federal AI use case data
- Reviewed national research and standards (e.g., NIST AI Risk Management Framework)
- Reviewed other states' AI policies and legislation

NIST = National Institute of Standards and Technology

In brief

- One-third of state agencies are using or piloting function-specific AI tools, and most employees at VITA-supported agencies use general-purpose AI tools.
- AI requirements for executive branch agencies lack risk-based criteria, clarity, and timely updates, undermining efficient and effective oversight.
- Policy and training gaps on general-purpose AI tools increase the state's risk of sensitive data exposure and other problems.

In brief (cont.)

- Near-term opportunities to use AI include modernizing dated state IT systems and reducing SNAP payment errors.
- Data readiness problems and poor coordination within VITA are among key barriers to agency AI adoption.
- Legal, technological, and policy issues can make it challenging to develop effective AI legislation, and systematic data on realized harms is limited.
- Specialized technology law expertise at DLS and routine, updated information on AI would support legislators' efforts to address AI-related harms.

SNAP = Supplemental Nutrition Assistance Program; DLS = Division of Legislative Services

In this presentation

Section 1	Background
Section 2	State agency use of function-specific AI
Section 3	State employee use of general-purpose AI
Section 4	Near-term opportunities & barriers
Section 5	Mitigating AI-related risks to the public

Section 1


Background

Modern “AI” generally refers to machine-based systems that infer how to make predictions, generate content, or achieve objectives

- No universally agreed-upon definition of “artificial intelligence”
- Modern AI definitions (e.g., OECD) generally distinguish AI systems from traditional software by their ability to *infer* how to achieve objectives rather than rely only on fixed, step-by-step instructions
- Advancements in “machine learning” techniques have driven recent progress in AI and adoption of AI tools
 - Machine-learning identifies patterns and improves performance through exposure to training data, rather than through human programming

OECD = Organisation for Economic Co-operation and Development

AI system capabilities can be generally classified as predictive, generative, or agentic




Predictive AI

Classifies, predicts, scores, or recommends.



Generative AI

Creates content such as text, images, code, or video.



Agentic AI

Plans and carries out multi-step tasks to achieve an objective.

Note: In practice, these systems can overlap. The list of capabilities for each type is illustrative and not exhaustive.

Increasing AI capabilities presents opportunities and risks for state government and the public

- Example opportunities:
 - improved citizen accessibility
 - improved service timeliness
 - improved healthcare diagnostics (e.g., medical imaging)
- Example risks:
 - decisions made based on biased or inaccurate information
 - exposure of children to harmful content
 - more effective fraud techniques and cybersecurity attacks

VITA provides centralized IT services to executive branch agencies and has established AI standards

- VITA provides IT infrastructure for 65 executive branch agencies and oversees agency IT procurement, projects, and security
 - State law requires executive branch agencies to follow VITA-established policies and standards for IT
- VITA plays a central role in the governance and risk management of AI for executive branch agencies
 - Approves use cases through its AI registry process
 - Articulates standards and policies for the development, procurement, and use of AI systems

Section 2

State agency use of function-specific AI


Function-specific AI tools perform specific agency operational or administrative tasks

- Examples include constituent-facing chatbots, fraud detection tools, tools that digitize and classify documents, and traffic incident detection systems
- Tools may be developed by agency staff, custom-built by vendors, acquired off the shelf, or embedded in commercial software
- VITA's AI Utilization Policy allows AI use only when it produces a “clear, positive outcome for Virginia's citizens” and is the optimal solution for the intended purpose

Finding – Function-specific AI use


Most agencies are not yet using function-specific AI, though many are considering it.

One-third of executive branch and independent agencies report using or piloting at least one function-specific AI tool



Source: JLARC survey on AI use at executive branch and independent state agencies (June–August 2026)

Agencies most frequently report using or considering function-specific AI to improve internal efficiency



Source: JLARC survey on AI use at executive branch and independent state agencies (June–August 2026)

State agencies reported using or piloting various function-specific AI tools (examples)

- Identify potential security threats in bags (DOC, deployed)
- Integrate and streamline various fiscal processes, including budget modifications, invoice processing, and financial reconciliation (VDEM, deployed)
- Analyze data across systems to improve project cost estimates (VDOT, pilot)
- Review permit applications for completeness (DEQ, pilot)
- Tailor resumes and cover letters for job seekers for specific job descriptions (Virginia Works, pilot)

NOTE: As of Sept. 2026. Examples only; not comprehensive. DOC = Department of Corrections; VDOT = VA Department of Transportation; VDEM = VA Department of Emergency Management; DEQ = Virginia Department of Environmental Quality

Finding – Uniform AI requirements

VITA's requirements for AI apply uniformly regardless of risk, undermining efficient and effective state oversight.

AI risks vary widely depending on their purpose, context, and data sensitivity and access

- National and international AI risk management guidance (e.g., NIST) recommends organizations prioritize AI risk management policies based on varying risks
- At least eight other states (e.g., Maryland, Washington, Texas) differentiate some of their AI requirements based on risk (e.g., low, medium, high)

NIST = National Institute of Standards and Technology

Most VITA AI requirements apply uniformly, regardless of risk level; VITA has not issued risk assessment guidance

- VITA standards contain risk mitigation measures that align with widely cited guidance, but most requirements are not proportionate to risk
- Almost all AI systems and tools require
 - approval from senior IT agency staff, specific VITA divisions, the VITA CIO, and the secretary of administration
 - internal and external validation of system design, data collection, and measurements in a system's deployment context
 - ongoing evaluation, including testing for degradation in model performance, risk assessment, incident tracking, and assessments for trustworthiness
- VITA has not provided guidelines to agencies for assessing AI risk

Uniform requirements for AI tools can be problematic

- For lower-risk uses, uniform requirements can delay or discourage adoption and create unnecessary burdens on VITA and agency staff
- For higher-risk uses, uniform requirements may not impose sufficiently stringent safeguards
 - e.g., some use cases may warrant review by OAG and relevant secretary, more frequent oversight
- Other states' approaches provide insights into how to classify AI risk

OAG = Office of the Attorney General

Recommendation #1

VITA should develop and publish an AI risk assessment instrument that provides agencies with clear, standardized criteria for evaluating potential AI use cases and determining their risk level.

Recommendation #2

VITA should develop and implement a risk-tiered approach for governing executive branch agencies' use of AI. As part of this approach, VITA should (i) allow agencies to approve low-risk AI use cases while still submitting required information to the AI registry, (ii) define agency and VITA responsibilities for oversight of low-, medium-, and high-risk use cases, and (iii) update, as appropriate, the AI Utilization Policy and the Enterprise Architecture Standard for AI to reflect tiered requirements.

Finding – Unclear AI standards

VITA's unclear AI standards make compliance challenging for agencies.

Agencies report that unclear state AI standards have made adoption of AI challenging

- Nearly half (48 percent) of executive branch agencies using function-specific AI cited unclear AI requirements as a major or moderate constraint
- Relatively few agencies found AI registry requirements “very clear”
 - which use cases must be submitted (30 percent), what information to submit (32 percent), VITA’s approval criteria (21 percent)
- JLARC staff’s review of the AI standards validates agency concerns
- Unclear requirements may lead agencies to interpret compliance differently, undermining state’s AI governance efforts

Recommendation #3

VITA should (i) revise the Enterprise Architecture Standard for AI and the AI Utilization Policy to ensure they are clear and actionable, (ii) solicit feedback from agencies to determine where further clarification is needed, and (iii) publish implementation guidance on its website.

Finding – Updating standards for emerging risks

State's AI standards have not kept pace with emerging risks, particularly those associated with agentic AI.

AI capabilities are advancing rapidly, and state's AI standards have not kept pace with evolving risks

- Agentic AI capabilities have improved substantially over past two years, are increasingly accessible, and pose heightened risk
 - Can independently execute multi-step workflows, creating risk of unintended or unauthorized actions (e.g., modifying, deleting, or sharing files)
 - Access controls need to be carefully reviewed and configured
- VITA has not updated its AI standards and policies frequently enough to keep up with AI's rapid evolution
 - VITA's Enterprise Architecture Standard for AI was last updated in November 2023 and does not refer to agents or related safeguards
- VITA staff acknowledge need to update standards and policies

Recommendation #4

VITA should (i) conduct a one-time targeted review to update the Enterprise Architecture Standard for AI to address emerging AI risks, including those related to agentic AI and data access controls, and (ii) develop and implement a process to annually review, and update as needed, all AI-related standards, policies, and guidance.

Section 3

State employee use of general-purpose AI

General-purpose AI tools can be used for a wide range of tasks

- General-purpose AI tools respond to user prompts or instructions by generating content (e.g., text, images, and code), or taking actions
- Include chatbots (e.g., Copilot, Claude, ChatGPT) and domain-specific tools that support broad tasks (e.g., coding assistants)
- May improve employee productivity and work quality in certain applications
- Irresponsible use increases risks, such as undetected errors or bias and sensitive data exposure

Findings – General-purpose AI use and risks

Data indicates that most employees at VITA-supported agencies use general-purpose AI tools.

Gaps in policies, guidance, and training for general-purpose AI tools increase state's risk of sensitive data exposure and other problems.

Most state agencies allow employees to use general-purpose tools, and most employees use them

- Almost all (97 percent) of surveyed executive branch and independent agencies allow staff to use general-purpose AI tools to some extent
- An estimated 76 percent of employees at VITA-supported agencies use a general-purpose AI tool (e.g., Copilot, ChatGPT, Claude) to some extent
- Since August 2025, VITA has offered an enterprise version of Copilot specifically for state agencies, but state employees also frequently use publicly available general-purpose AI tools

State is operating at a higher-than-necessary risk level from potential employee misuse of general-purpose AI tools

- Lack of statewide acceptable use policy that articulates responsible use expectations
 - Makes holding employees accountable for misuse challenging
 - As of June 2026, only 28% of agencies had written staff policies or guidance; 90% said a statewide policy would help support responsible use
- Unclear guidance on which general-purpose AI tools are approved for employee use, increasing risks of accidental misuse of public tools
- Limited staff training on responsible use and limitations of AI
 - 92% said general training for all employees would help support responsible use

Recommendation #6

VITA should (i) develop an acceptable use policy governing executive branch employees' use of general-purpose AI tools, (ii) review and update the policy annually as needed, (iii) require executive branch agencies to adopt VITA's policy or develop their own policy that includes the requirements established by VITA's policy, and (iv) require executive branch agencies to maintain documentation of employees' acceptance of their agency's policy as part of initial onboarding and annually thereafter.

Recommendation #7

VITA should (i) post prominently on its website clear information on which general-purpose AI tools are approved for state employee use, and (ii) establish a process to ensure this information is updated promptly when new tools are approved or their approval status changes.

Recommendation #8

The General Assembly may wish to consider amending the Code of Virginia to require (i) the chief information officer of the Commonwealth to develop and annually update a curriculum and materials for state employee AI literacy training, including responsible use of general-purpose AI tools; and (ii) that all state agencies certify annual employee completion.

Section 4

Near-term opportunities & barriers

Federal agencies and other states are increasingly using AI tools for government operations and services

- AI models now perform at or above human baseline levels on tasks such as language understanding, reading comprehension, and software engineering*
- Federal agencies' AI use cases grew from 6 in 2016 to 340 in 2025, according to OMB
- “Surge” in AI use cases in recent years, according to NASCIO’s 2025 survey of state chief information officers

*Source: Stanford University

OMB = U.S. Office of Management and Budget; NASCIO = National Association of State Chief Information Officers

Finding – Near-term opportunities

Near-term potential opportunities to leverage AI tools include modernizing state IT systems and reducing SNAP payment errors.

AI coding assistants could help modernize state IT applications

- As of August 2025, VITA had identified 865 Commonwealth IT applications needing modernization, including 344 requiring urgent attention
- Several state agencies are already using or considering AI tools to help modernize their IT systems
- In 2025, Texas passed a law directing its IT agency to help agencies modernize or replace legacy IT systems with AI assistance

AI tools could support state's effort to reduce SNAP payment error rate and associated federal cost

- Federal changes to SNAP under H.R. 1 (2025) require states to contribute to SNAP benefit costs if payment error is 6 percent or higher
- Reducing payment error rates could create significant cost savings for state
 - VDSS estimated state would owe \$270 million annually beginning in FY28 if errors remained at FY24 rate
- Several other states are either using or planning to use AI technologies to help reduce SNAP payment errors
- As of September 2026, VDSS had deployed two tools to detect or prevent SNAP payment errors, identified other potential use cases

VDSS = Virginia Department of Social Services

Recommendation #9

VITA should (i) establish an expedited review and approval determination process for AI coding assistance tools proposed by state agencies for IT system modernization; (ii) evaluate the feasibility of establishing a statewide contract for one or more AI coding assistance tools; and (iii) develop specific technical and risk-mitigation guidance and requirements for state agencies' use of AI coding assistance tools.

Recommendation #10

VITA should prioritize its review and approval determination of AI use cases identified by the Virginia Department of Social Services to reduce Supplemental Nutrition Assistance Program payment errors.

Finding – Coordination and efficiency challenges within VITA

Coordination and efficiency challenges within VITA will complicate AI adoption efforts if not addressed.

Agencies report frustration with cumbersome and uncoordinated VITA review and approval processes

- VITA organized into seven major functional areas, most of which play a role in the governance of AI systems and tools
 - AI projects can require as many as 20 separate reviews, including multiple reviews by the same division
- Agencies reported problems due to poor coordination at VITA, including conflicting information, unclear processes, and need to use personal relationships to expedite projects
- Delays and coordination problems in VITA's AI governance part of broader organizational problems that have persisted for years
 - Example: Only 29 percent of VITA staff reported VITA's work processes were highly efficient in 2025 internal culture survey

Recommendation #5

VITA should develop and publish a comprehensive process map detailing the required review and approval steps for state agency AI projects and procurements.

Recommendation #12

VITA should (i) identify where agencies experience poor coordination, confusion, and delays when completing VITA's required processes to adopt artificial intelligence tools, and (ii) determine and address the root causes of these problems.

Lack of identified need, data readiness, and funding also cited as barriers to state agency adoption of AI tools

- Most agencies not pursuing AI (72 percent) reported they had not identified a business need or clear use case
 - Some may lack a business case (e.g., cost, complexity, risk), but others may be unaware of how AI could help their operations
- Majority of agencies using or pursuing function-specific AI tools cited data deficiencies as major or moderate constraint
 - Poor data (e.g., inaccurate, incomplete, unrepresentative) increases risk of poor outputs
- More than two-thirds of agencies pursuing or using AI tools cite funding constraints

Section 5

Mitigating AI-related risks to the public

Public and businesses are increasingly using AI technologies

- About half of U.S. adults (49 percent) used AI chatbots in 2026, up from one-third in 2024, according to Pew Research Center surveys
 - Reported uses include information searches, medical advice, and emotional support
- 64 percent of teens report using AI chatbots, and 28 percent use them at least daily (Pew, 2025)
- About 23 percent of Virginia businesses reported using AI as of late July 2026 (U.S. Census Bureau)

Wide range of potential AI-related harms exists, and little systematic data is available on realized harms

- Experts indicate AI tools can cause or contribute to public harm when misused or used irresponsibly
 - MIT AI Risk Repository has documented over 1,700 risks (e.g., toxic content, discrimination, cyber weapons)
- Limited systematic data exists on realized harms, but evidence indicates AI is contributing to harmful activity
 - FBI's 2025 Internet Crime Report: more than 22,000 AI-related complaints nationally; reported losses of ~\$22.8 million in Virginia
 - Press reports about recent AI-related cybersecurity incidents
- States are increasingly enacting AI legislation, with mixed success

Finding – Challenges developing legislation to address AI harms

Developing effective AI legislation to address potential AI harms can be challenging and require in-depth technology law and policy expertise.

Challenges: existing law & constitutional limits

- Many Virginia laws are technology neutral and may already address some AI-related harms (e.g., consumer deception, discriminatory credit decisions, identity theft)
 - Uncertainty about their scope makes it difficult to know whether new legislation is needed
 - Detecting and proving harm from AI systems, such as discrimination, may be difficult even where existing law applies
- AI-related laws may conflict with constitutional protections, such as free speech
 - Example: Recent CA and HI election AI deepfake laws blocked or struck down

Challenges: unclear definitions and assigning liability

- No universally accepted definition of AI, which can result in unclear or overly broad legislative definitions
 - Broad definition in previous Colorado AI law led to disagreement over which technologies it covered (law later repealed and replaced)
- Multiple parties (developers, deployers, end users) can contribute to AI harms, making liability difficult to assign
 - Agentic AI tools and publicly downloadable and modifiable models (e.g., open-weight models) further complicate liability

Note: With open-weight models, a developer releases the model's trained parameters (its "weights") for public download and modification. Anyone who downloads an open-weight model can adjust the model's weights for their own purposes.

Challenges: compliance feasibility, context dependence, and limited harms data

- Some compliance requirements may be difficult or costly to meet, such as explaining how advanced AI systems reach decisions or reliably detecting AI-generated content
- Same AI tool can be beneficial or harmful, so broad restrictions can also limit beneficial uses
- Systematic data on realized harms is very limited, making both over- and under-regulation possible
 - Three states (CA, NY, IL) have enacted legislation requiring frontier AI developers to report critical safety incidents, which should increase information on some AI-related harms

Finding – Technology law expertise at DLS

Limited technology law expertise within DLS may challenge its ability to provide effective guidance on AI-related legislation.

DLS serves a key role to ensure legislation achieves its intended purposes, avoids unintended consequences

- Statutorily tasked with drafting bills and advising legislators on constitutionality or probable legal effect of proposed legislation upon request
- Required to prepare summaries of existing laws affected by proposed legislation, upon request of any member

Limited technology law expertise within DLS may challenge its ability to provide effective guidance on AI-related legislation

- Advising on AI legislation requires specialized technology law expertise
 - e.g., assigning liability across developers, deployers, and users, or judging whether explainability requirements are technically feasible
- As of September 2026, DLS does not have on staff, or access to, an attorney specializing in technology law
 - Legislators and DLS staff indicated need for technology law expertise
- Obtaining ongoing expertise could include hiring, partnering with a university, or retaining outside counsel

Recommendation #13

The General Assembly may wish to consider including funding in the Appropriation Act for the Division of Legislative Services to obtain, on an ongoing basis, specialized technology law expertise needed to effectively advise legislators on AI legislation.

Finding – Up-to-date information on AI

Up-to-date information on AI-related harms could help the General Assembly navigate legislative challenges.

Up-to-date information on AI-related harms could help General Assembly navigate legislative challenges

- Systematic and defensible data on AI-related harms is limited, and up-to-date information on AI-related developments is needed
- JCOTS is statutorily charged with advising Virginia leaders on science and technology matters and could keep legislators informed on key AI trends
- A short, annually updated legislator's guide could cover key AI terminology, legislative considerations, benefits and harms, and recently enacted legislation
 - Similar to JLARC's guide on VRS

Recommendation #14

The General Assembly may wish to consider directing the Joint Commission on Technology and Science to publish and annually update an informational legislator's guide on artificial intelligence for General Assembly members and make the guide available no later than October 1 of each year.

JLARC staff for this report

Kimberly Sarte, Associate Director for Ongoing Oversight & Fiscal Analysis

Drew Dickinson, Chief Legislative Analyst

Laura White, PhD, Senior Legislative Analyst

Madison McCaffrey, Senior Associate Legislative Analyst

Alexandra Jansson, Principal Legislative Analyst for Ongoing Oversight