

Summary: Governance and Oversight of Artificial Intelligence in Virginia

WHAT WE FOUND

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

Function-specific AI tools are designed to perform specific agency operational or administrative tasks (e.g., a chatbot that answers routine public inquiries, a fraud detection system, and a meeting transcription service), rather than being adaptable across a broad range of uses like general-purpose AI tools (e.g., ChatGPT, Claude, Copilot Chat). As of June 2026, about a third of agencies reported using or piloting at least one function-specific AI tool. However, an additional 30 percent of agencies were actively considering or pursuing adoption of a function-specific AI tool.

VITA's requirements for AI use are not proportionate to risk, undermining efficient and effective state oversight

Most of VITA's AI standards and requirements apply uniformly regardless of risk, even though AI risks vary widely in type and severity depending on their purpose, deployment context, data sensitivity, and access to organizational files and systems. National and international guidelines recommend that organizations prioritize AI risk management policies and resources based on each AI tool's risk likelihood and severity, and at least eight other states differentiate some of their AI requirements based on risk. Largely uniform requirements can delay or discourage the adoption of low-risk use cases and create unnecessary burdens on VITA and agency staff. Likewise, uniform requirements may not impose sufficiently stringent safeguards on high-risk use cases.

In addition to uniform requirements, VITA has not provided agencies with guidelines for how to assess the risk of AI tools and uses, even though it requires agencies to identify risks as part of the overall AI approval process.

WHY WE DID THIS STUDY

In 2025, the Joint Legislative Audit and Review Commission directed staff to review the use, governance, regulation, and oversight of artificial intelligence in the Commonwealth.

ABOUT ARTIFICIAL INTELLIGENCE

Artificial intelligence (AI) generally refers to machine-based systems that infer how to make predictions, generate content, or achieve objectives, tasks commonly associated with human intelligence. Advancements in machine learning have driven recent progress in AI and the adoption of AI tools. Machine learning enables computer systems to identify patterns and improve performance through training data rather than explicit human programming. Modern AI systems create both opportunities, such as improved operational efficiency, and risks, such as biased recommendations and decisions, for state government and the public. Responsibility for the governance and oversight of AI in Virginia is spread across various entities. Within state government, the Virginia Information Technologies Agency (VITA) has established standards governing how state agencies should adopt and use AI. AI oversight also involves the General Assembly, its support agencies, and existing regulatory and enforcement agencies and authorities.

Unclear AI standards make compliance challenging for agencies, and standards have not kept pace with emerging risks

Agencies report that unclear state AI standards have made adoption of AI challenging. About half (48 percent) of executive branch agencies using function-specific AI cited unclear requirements as a major or moderate constraint on their pursuit or implementation of AI. JLARC staff's review of state AI standards found multiple provisions are unclear about requirements and how agencies should demonstrate compliance. Unclear standards may lead agencies to interpret compliance differently, undermining the state's efforts to ensure AI systems are safe and effective.

In addition to being unclear, VITA's standards have not sufficiently kept pace with emerging risks. For example, agentic AI capabilities have improved dramatically in the last two years, and AI agents can autonomously plan and take actions to achieve a specific goal without needing a human to define each step. However, VITA's standards do not refer to agents or necessary safeguards for autonomous tools, including specific controls needed to help mitigate risks associated with agents' access to state systems and data. These governance gaps exist, in part, because VITA has not updated its AI standards frequently enough to reflect AI's rapid technological advancements.

Most employees at VITA-supported agencies use general-purpose AI tools, and their undirected use can increase data exposure risk

Almost all executive branch and independent state agencies (97 percent of agencies surveyed by JLARC) allow staff to use general-purpose AI tools, though many restrict their use to some extent. Among state employees served by VITA, an estimated 76 percent used a general-purpose AI tool to some extent during a one-month period. General-purpose AI tools include chatbots, such as Copilot Chat, Claude, and ChatGPT, as well as tools tailored for specific domains, such as coding assistants (e.g., Claude Code, OpenAI Codex, GitHub Copilot). VITA and agencies report that staff use these tools for a variety of purposes, including improving code, summarizing documents, and drafting emails.

In August 2025, VITA started providing an enterprise version of Copilot to all supported agencies, but usage data indicates state employees frequently use publicly available general-purpose AI tools, such as the public versions of Copilot, ChatGPT, Gemini, and Claude. Employer-provided enterprise AI tools are generally more secure than publicly available tools because they can offer enhanced data protections (e.g., data is not used to train the AI model), greater administrative oversight, and additional configuration options (e.g., the ability to disable features). In contrast, using publicly available general-purpose AI tools carries greater risk of data exposure if employees do not use them responsibly.

Virginia is currently exposed to higher-than-necessary risk of state employee misuse of publicly available AI tools because it lacks a statewide acceptable use policy, staff

training on responsible AI use, and clear guidance on which specific tools are authorized for particular functions.

Lack of state employee guidance and training on general-purpose AI tools increases risk of improper use and may hinder accountability

Virginia has little formal statewide guidance for state employees on the use of general-purpose AI. In addition, VITA does not provide clear guidance on which tools are approved for use by agency staff. Acceptable use policies for general-purpose AI tools can help organizations reduce risks associated with these tools, such as employees unintentionally exposing sensitive information. Ninety percent of agencies surveyed indicated that a statewide acceptable use policy for general-purpose AI tools, with the ability for agencies to establish additional restrictions, would be either very helpful or somewhat helpful in supporting effective and responsible use by staff.

State IT leaders and VITA staff also identified insufficient AI literacy among agency staff as a primary barrier to their use of AI. Overall, 76 percent of agencies that currently use or are actively pursuing AI reported that insufficient staff AI literacy was a major or moderate constraint on their agency's ability to pursue or use AI tools or systems. Most agencies indicated that general training for state employees on the effective and responsible use of general-purpose AI tools would be helpful.

Agencies indicate staff training would help support effective and responsible general-purpose AI use



SOURCE: JLARC survey on artificial intelligence use at state agencies (June–August 2026).

NOTE: N = 79 executive branch and independent agencies. Figures do not sum to 100 due to rounding and because one agency reported "I don't know."

Near-term opportunities exist for greater use of AI in Virginia state government, but agencies also face barriers

State agencies are already using or pursuing a wide range of AI use cases to streamline operations; reduce costs; improve staff productivity, data quality, and accessibility; and assist with compliance, oversight, and fraud detection. Two significant near-term opportunities for additional state AI use include:

- **Modernizing the state's dated IT systems**—As of August 2025, VITA had identified more than 850 state IT applications in need of modernization, in-

cluding over 300 requiring urgent attention. AI coding assistants, if implemented with appropriate controls, could help state efforts to update outdated software and modernize legacy IT systems; and

- **Reducing Supplemental Nutrition Assistance Program (SNAP) payment error rates**—Federal changes to SNAP under H.R. 1 (2025) require states to contribute to the cost of SNAP benefits if their payment error rate is at or above 6 percent. In January 2026, using the FY24 payment error rate (11.50 percent), the Virginia Department of Social Services (VDSS) estimated that Virginia would be required to pay \$270 million annually beginning in FY28, assuming the error rate remained the same. Several states are either using or planning to use AI technologies to reduce SNAP payment errors, and the federal government has encouraged states to pursue advanced technologies, including AI tools, to improve SNAP program administration. As of September 2026, VDSS had deployed two AI-enabled tools to help prevent and detect SNAP errors and had identified other potential AI use cases that could help further reduce errors.

While many agencies are interested in further exploring and pursuing AI tools, state agency staff cited several common barriers to their ability to adopt new AI tools, including a lack of awareness of potential business use cases, data readiness challenges, lack of coordination of the review and approval processes within VITA, and funding constraints.

Multiple issues can make AI challenging to regulate through legislation

AI tools can be both beneficial and harmful to the public depending on how they are used. As a result, the General Assembly and legislatures in other states have been considering and passing legislation to protect the public from potential AI-related harms. Developing effective AI legislation that balances these potential risks and benefits can be difficult because of the rapid pace of AI development, as well as various legal, technological, and policy complexities.

Existing laws and regulations may already address harms—Some existing state laws and regulations may already cover AI-related harms, but uncertainty about their scope makes it difficult for legislators to know whether new legislation is needed. For example, many Virginia laws are technology-neutral (e.g., those pertaining to consumer deception) and regulate the *underlying conduct, practices, or decisions*, rather than the *means* by which they occurred (e.g., a human or an automated system). However, even though existing laws and regulations may address some AI-related harms, legal gaps and evidentiary challenges may remain.

Potential constitutional conflicts—Proposed AI legislation and regulations could potentially conflict with federal or state constitutional protections, such as rights to

free speech, privacy, and due process. State AI laws could also conflict with other constitutional provisions, such as the U.S. Constitution’s Supremacy Clause (which articulates that federal laws take precedence over state laws) and Commerce Clause (which gives Congress the power to regulate interstate commerce). Some AI-related legislation in other states has already been found to conflict with constitutional protections.

Unclear technical definitions—There is no universally accepted definition of AI, which can result in legislative definitions that are unclear or overly broad. For example, an overly broad definition of “AI systems” in a 2024 Colorado law contributed to disagreement over which technologies the law covered. Colorado eventually repealed and replaced the law, adopting a narrower and more specific definition of which technology would be regulated.

Assigning liability for AI-related harms—The development, deployment, and use of AI systems often involves multiple parties, making it difficult to assign responsibility for any resulting harms. Harms may also be caused, in part, by design choices or uses unknown to other parties. If legislation does not clearly delineate responsibility for potential harms, injured parties will struggle to seek redress, ultimately undermining the law’s effectiveness.

Difficult compliance requirements—Current technological limitations may make some compliance requirements in AI-related legislation difficult to fully meet. For example, laws requiring deployers to disclose the principal reasons for an adverse consequential decision made or substantially influenced by an AI system may be difficult to comply with, because the internal decision-making of many advanced AI systems is often not fully explainable. Some states have instead pursued other measures, such as requiring frontier AI developers to report critical safety incidents.

Risks depend on context—AI tools are often “dual use,” meaning the underlying technology could be used in both beneficial and harmful ways. Additionally, the same AI tool can have different risk levels depending on its application and context. For example, a general-purpose chatbot presents minimal risk when being used to help draft a meeting agenda but can create significant risk when providing medical advice.

Limited data on realized harms—Systematic data about harms enabled or caused by AI systems is currently very limited, making it difficult to gauge their breadth, prevalence, and severity. Without this insight, both over- and under-regulation of AI is possible.

Technology law expertise and up-to-date information on AI would help legislators effectively address AI-related harms

To help legislators navigate the challenges posed by AI-related legislation, the Division of Legislative Services (DLS) needs staff with specialized technology law expertise, particularly related to AI, or access to such expertise. In interviews, both legislators

and DLS staff indicated the need for this specialized knowledge. However, as of September 2026, DLS does not have on staff, or access to, an attorney specializing in technology law to advise the General Assembly.

Access to technology law expertise could also help legislators determine how AI-related legislation is referred and reviewed during the legislative session, as needed. Proposed legislation may often raise both policy and technology issues, requiring a determination of whether it should be assigned to a policy committee or a technology committee. Similarly, technology law expertise can help legislators determine when a bill requires in-depth review by a legislative commission. The Joint Commission on Technology and Science (JCOTS) may be best suited to address legislation seeking to regulate the *development or deployment* of AI technologies, but other commissions may be better able to evaluate legislation seeking to regulate the *use* of AI technologies in a particular policy context.

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

Legislators will need to evaluate AI-related legislation for the foreseeable future, but comprehensive information on realized AI-related harms is limited, making it difficult to gauge their prevalence and the need for regulation. An up-to-date AI guide for legislators could be helpful as they develop or deliberate AI legislation. Such a guide could cover topics like important AI terminology; key legislative considerations; potential, actualized, and emerging benefits and harms from AI use; summaries of recently enacted AI legislation in Virginia and elsewhere; and recommendations of reputable resources for further information about AI.

WHAT WE RECOMMEND

The following recommendations include only those highlighted in the report summary. The complete list of recommendations is available on page ix.

Legislative action

- Require the chief information officer of the Commonwealth to develop and annually update a curriculum for state employee AI literacy training, including responsible use of general-purpose AI tools, and require state agencies to certify annual employee completion.
- Include funding in the Appropriation Act for DLS to obtain, on an ongoing basis, specialized technology law expertise.
- Direct JCOTS to provide an informational legislator’s guide on AI for General Assembly members by October 1 of each year.

Executive action

- VITA should publish an AI risk assessment instrument for agencies and implement a risk-tiered approach for governing executive branch agencies' use of AI.
- VITA should ensure its standards and policies for AI systems and uses are
 - clear and actionable, and
 - updated to address emerging risks, such as those related to agentic AI, and reviewed annually for any further updates needed.
- VITA should develop an acceptable use policy governing executive branch employees' use of general-purpose AI tools and clearly post on its website information on which general-purpose AI tools are approved for state employee use.
- VITA should establish an expedited review and approval determination process for agencies' proposals to use AI coding assistance tools for IT system modernization, and develop technical and risk-mitigation guidance and requirements for their implementation.
- VITA should prioritize its review and approval determination of AI use cases identified by VDSS to reduce SNAP payment errors.
- VITA should identify where agencies experience poor coordination, confusion, and delays when completing VITA's required processes to adopt AI tools, and determine and address the root causes of these problems.

