

Report to the Governor and the General Assembly of Virginia

Governance and Oversight of Artificial Intelligence in Virginia

2026



COMMISSION DRAFT

Joint Legislative Audit and Review Commission

Chair

Senator Mamie E. Locke

Vice-Chair

Delegate Betsy B. Carr

Senator R. Creigh Deeds
Delegate Charniele L. Herring
Senator Ryan T. McDougle
Senator Jeremy S. McPike
Senator Mark Obenshain
Delegate Cia Price
Delegate Sam Rasoul
Delegate David Reid
Delegate Marcus B. Simon
Delegate Josh Thomas
Delegate R. Lee Ware
Delegate Tony O. Wilt

Staci Henshaw, Auditor of Public Accounts

Director

Hal E. Greer

JLARC staff for this report

Kimberly Sarte, Associate Director
Drew Dickinson, Chief Legislative Analyst, Project Leader
Laura White, Senior Legislative Analyst
Madison McCaffrey, Senior Associate Legislative Analyst
Alexandra Jansson, Principal Legislative Analyst for Ongoing Oversight

Managing Editor: Jessica Sabbath
Information graphics: Nathan Skreslet

Contents

Summary	i
Recommendations	ix
Chapters	
1. Artificial Intelligence in Virginia	1
2. Function-Specific AI Use by State Agencies	9
3. General-Purpose AI Use by State Employees	23
4. Near-Term Opportunities and Barriers	31
5. Mitigating AI-Related Risks to the Public	41
Appendixes	
A. Study resolution	55
B. Research activities and methods	57
C. Agency responses	63
D. Key factors affecting costs and availability of AI	68
E. Recent AI legislation passed by the Virginia General Assembly	74
F. Risk-based approaches to AI oversight in other states	77

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.

Recommendations: Governance and Oversight of Artificial Intelligence in Virginia

RECOMMENDATION 1

The Virginia Information Technologies Agency should develop and publish an artificial intelligence (AI) risk assessment instrument that provides agencies with clear, standardized criteria for evaluating potential AI use cases and determining their risk level. (Chapter 2)

RECOMMENDATION 2

The Virginia Information Technologies Agency (VITA) should develop and implement a risk-tiered approach for governing executive branch agencies' use of artificial intelligence (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 Artificial Intelligence Utilization Policy and the Enterprise Architecture Standard for Artificial Intelligence to reflect tiered requirements. (Chapter 2)

RECOMMENDATION 3

The Virginia Information Technologies Agency should (i) revise the Enterprise Architecture Standard for Artificial Intelligence and the Artificial Intelligence 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. (Chapter 2)

RECOMMENDATION 4

The Virginia Information Technologies Agency should (i) conduct a one-time targeted review to update the Enterprise Architecture Standard for Artificial Intelligence 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. (Chapter 2)

RECOMMENDATION 5

The Virginia Information Technologies Agency should develop and publish a comprehensive process map detailing the required review and approval steps for state agency artificial intelligence projects and procurements. (Chapter 2)

RECOMMENDATION 6

The Virginia Information Technologies Agency (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 the 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 during initial onboarding and annually thereafter. (Chapter 3)

RECOMMENDATION 7

The Virginia Information Technologies Agency 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. (Chapter 3)

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. (Chapter 3)

RECOMMENDATION 9

The Virginia Information Technologies Agency should (i) establish an expedited review and approval determination process for artificial intelligence coding assistance tools proposed by state agencies for IT system modernization; (ii) evaluate the feasibility of establishing a statewide contract for one or more artificial intelligence coding assistance tools; and (iii) develop specific technical and risk-mitigation guidance and requirements for state agencies' use of artificial intelligence coding assistance tools. (Chapter 4)

RECOMMENDATION 10

The Virginia Information Technologies Agency should prioritize its review and approval determination of artificial intelligence use cases identified by the Virginia Department of Social Services to reduce Supplemental Nutrition Assistance Program payment errors. (Chapter 4)

RECOMMENDATION 11

The Virginia Information Technologies Agency should make available to state agencies a regularly updated list of piloted and deployed state agency artificial intelligence (AI) use cases to increase awareness of potential AI applications in state government that can support agency operations and service delivery. (Chapter 4)

RECOMMENDATION 12

The Virginia Information Technologies Agency (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. (Chapter 4)

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 artificial intelligence legislation. (Chapter 5)

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. (Chapter 5)

1 Artificial Intelligence in Virginia

In December 2025, the Joint Legislative Audit and Review Commission (JLARC) directed its staff to review the use, governance, regulation, and oversight of artificial intelligence (AI) in the Commonwealth. JLARC staff were directed to review AI use in Virginia state government and its opportunities and risks; whether agency staff have appropriate skills to use AI effectively; and whether the state’s AI policies and processes enable innovation while mitigating risks. Staff were directed to identify approaches that would help the General Assembly identify the most substantial risks to the public from AI use outside state government and address them through legislation, where necessary. Staff were also directed to identify key factors that affect the cost and availability of AI systems. (See Appendix A for the study resolution and Appendix D for a discussion of factors affecting AI systems’ cost and availability.)

JLARC staff used various methods to address the study mandate, including interviews with state agency staff, national subject matter experts, and industry stakeholders; a statewide agency survey; and reviews of other states’ approaches to using and regulating AI. JLARC staff interviewed leaders and staff at the Virginia Information Technologies Agency (VITA) and state agencies of various sizes and functions, such as the Virginia Department of Motor Vehicles, the Virginia Department of Social Services, and the State Corporation Commission; national experts on AI technologies and policy; legislative staff in Virginia and in other states; and industry stakeholders. Staff also conducted a statewide survey of 80 executive branch and independent agencies, and reviewed AI legislation in other states; VITA policies, standards, and documents; and national AI risk management standards. (See Appendix B for a detailed description of research methods.)

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

There is no universally agreed-upon definition of “artificial intelligence,” but AI systems are generally understood as machine-based systems that perform tasks that are commonly associated with human intelligence. This includes tasks such as learning, reasoning, solving problems, perceiving information, using language, and operating under uncertain or changing conditions.

Modern AI definitions 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, typically defined by human programmers. The Organisation for Economic Co-operation and Development (OECD) defines an “artificial intelligence system” as

a machine-based system that, for explicit or implicit objectives, infers, from the input it receives, how to generate outputs such as predictions, content, recommendations, or decisions that can influence physical or virtual environments. Different AI systems vary in their levels of autonomy and adaptiveness after deployment.

Even this widely used definition, however, may be too broad to be used alone in regulations and governance because it covers a wide range of systems of varying impacts and risks (e.g., deciding public benefits eligibility versus generating a draft agenda). The OECD notes that additional criteria may be needed to narrow or tailor the definition depending on the context.

Advancements in machine learning techniques have driven recent progress in AI and adoption of AI tools

“Expert systems” are another example of AI.

These systems are programmed based on expert knowledge or criteria and produce outputs consistent with that programming, rather than inferring rules from data. Examples include tax preparation and logistics scheduling software.

Although AI comprises a variety of techniques and systems, advancements in machine learning have driven much of the field’s progress in recent years (Figure 1-1) (sidebar). Generally, machine learning involves techniques that enable computer systems to identify patterns and improve their performance through exposure to training data, rather than through explicit human programming.


Some of the most advanced AI systems are based on “deep learning” and rely on complex architectures called “deep neural networks.” Deep learning is a machine learning technique involving networks with many layers of interconnected units (or “nodes”) that identify patterns and build increasingly abstract numerical representations of input data, such as text or images. Many well-known modern generative AI tools, such as ChatGPT, Claude, Gemini, and Copilot Chat, are built on deep neural network architectures.

Flawed or biased outcomes can arise from other design and deployment choices, such as using poor proxy variables as the basis for predictions and using a model in a context for which it was not designed.

The complexity of deep neural networks contributes to both the capabilities of many modern AI systems and their risks, such as limited transparency and learned biases. Because these networks develop numerical representations of patterns in their training data rather than relying solely on explicit programmer instructions, they can generate outputs that would be impractical to program directly, such as highly personalized responses or recommendations. However, because of this process, even an AI system’s developers may be unable to fully explain how it reached a specific output, a limitation known as the “black box” problem. This limitation can make modern AI systems’ outputs difficult to verify, audit, or challenge. Additionally, because a deep neural network’s outputs reflect the data the system has learned from, poor or biased training data can lead to systematically flawed or biased outcomes (sidebar).

FIGURE 1-1

Machine learning is a subfield of the broader field of artificial intelligence, where systems learn and apply patterns from data



SOURCE: JLARC staff analysis.

Advances in machine learning have also improved natural language processing, which has made AI tools more accessible to the general public. Individuals without programming experience can interact with an AI system simply by entering a text prompt or speaking to it. Other factors contributing to the increasing capabilities and availability of machine-learning-based systems include the vast amounts of data online and advancements in hardware, such as graphics processing units, that can perform many tasks in parallel rather than sequentially.

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


For governance purposes, modern AI system capabilities can generally be understood as predictive, generative, or agentic (Figure 1-2). In practice, however, these categories can overlap. For example, a system (or combinations of systems) may use predictive models to identify potential cases of fraud and generative tools to produce a report on its findings.

Predictive AI systems learn from and analyze data to classify, score, rank, or forecast. These systems produce an assessment about specific cases rather than create content. Examples of predictive AI systems include medical image readers to identify a possible tumor, fraud detection systems, and credit scoring tools for default risk.

Generative AI systems create content, including text, images, audio, video, and computer code, based on patterns in the data used to train them. Large language models power widely available chatbots, such as ChatGPT, Claude, Gemini, and Copilot Chat, as well as other types of tools, such as writing, coding, and translation software. Examples of tasks performed by generative AI systems include drafting a summary of a long document, answering routine public inquiries through a chatbot, and generating realistic images, video, or voice recordings.

Agentic AI systems plan and carry out multi-step tasks to achieve an objective, with limited human involvement. Unlike a system that returns an output for a person to act on, an agentic system takes the actions itself, working through its connections to other software tools and data sources. Examples include a consumer tool that compares products, completes purchases, and schedules delivery, and a software coding tool that diagnoses bugs, writes fixes, and submits the code to a software engineer for review. Agentic systems vary widely in how much autonomy they have.

FIGURE 1-2
AI system capabilities can be grouped into three broad types



SOURCE: JLARC staff analysis.

NOTE: In practice, these systems can overlap. For example, agentic AI systems may use both predictive and generative AI systems in a workflow. Similarly, predictive AI systems (e.g., medical imagery diagnostics) may be paired with generative AI systems (e.g., summary report generator). The list of capabilities for each type is illustrative and not exhaustive.

Modern AI tools can assist with specific tasks or a broad range of tasks

For governance purposes, modern AI tools can also be classified by the breadth of tasks they are designed to perform: (1) function-specific tools built to perform a particular task (e.g., a chatbot that answers routine public inquiries, a fraud detection system, and a meeting transcription service), and (2) general-purpose tools (e.g., ChatGPT, Claude, Gemini, Copilot Chat) that are adaptable across a wide range of tasks. In practice, the line between function-specific and general-purpose tools is not always clear, such as with highly capable coding assistance tools.

The two broad types of AI tools are typically adopted and used differently, which has implications for how each is governed. Function-specific tools are selected, procured, and deployed by an organization for a defined purpose, and can therefore be reviewed and approved before they are used. Although organizations may formally deploy general-purpose AI tools, many of these tools are widely available to individual employees. Employees can decide whether and how to apply general-purpose AI tools to their day-to-day work, within any limits set by organizational policies. (See Chapter 2 for a

discussion on Virginia state agencies' use and governance of function-specific AI tools and Chapter 3 for a discussion on state employees' use and governance of general-purpose AI tools.)

Increasing AI capabilities present both opportunities and risks for state government and the public

Modern AI systems create both opportunities and risks for state government operations and public services. Examples of potential benefits include greater citizen accessibility, such as translating government websites into multiple languages; faster support for citizens, such as helping frontline staff locate information across large volumes of documents; and stronger cybersecurity defenses, such as identifying and responding to abnormal activity on state networks. Conversely, examples of potential risks for state government operations and services include providing inaccurate information to citizens, such as through a government chatbot; making poor or biased decisions on government benefits because of poor or biased data or insufficient human oversight; and exposing confidential or private data, such as through an agentic tool that has access to this data and sends it through an employee's email account. (See Chapter 4 for a discussion of near-term opportunities to use AI in Virginia state government and barriers to adoption and Chapters 2 and 3 for a discussion on AI-related risks and mitigation strategies in Virginia state government.)

Advances in machine-learning-based AI systems also present opportunities and risks for the broader public. Examples of potential opportunities include enhanced productivity, such as by automating routine tasks like meeting transcription; improved healthcare diagnostics, such as through advanced medical imaging techniques; and personalized learning, such as through conversational chatbots. Examples of risks to the public include biased decisions by businesses about access to employment, housing, or insurance; chatbots exposing harmful content to children; and the use of more effective fraud techniques by criminals. (See Chapter 5 for a discussion on AI-related harms to the public and strategies the General Assembly could consider to understand and mitigate them.)

Various executive and legislative branch agencies are involved in governance and regulation of AI

Responsibility for the governance and oversight of AI in Virginia is spread across various entities, which is partially a result of its wide range of use cases and its rapid development and adoption. Within state government, VITA provides IT infrastructure and services to most executive branch agencies and has established the standards governing how those agencies should adopt and use AI. Outside state government, mitigating AI-related harms to the public involves the General Assembly, its support agencies, and the regulatory and enforcement agencies and authorities that already oversee particular sectors and practices.

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

State law requires judicial, legislative, and independent agencies and authorities to follow VITA security standards. However, these entities, as well as certain higher education institutions, are not required to use VITA infrastructure and services.

VITA is Virginia’s consolidated information technology (IT) agency and is responsible for providing infrastructure and oversight of state agency IT functions. VITA provides infrastructure services, including computers, internet, phone, email, and security, for 65 executive branch state agencies, as well as enterprise-wide software applications. It also oversees agencies’ IT procurement, project management, security, and strategic planning (sidebar).

VITA has published AI standards that executive branch agencies must follow. VITA’s AI Utilization Policy includes core governing requirements such as the ethical use of AI; a citizen-benefit rule for its use; mandatory registration and approval by VITA; disclosure requirements for certain contexts; and data protections. As part of the approval process, VITA requires agencies to submit information about “any internal or external AI system that an agency seeks to develop, implement, or procure” into a central AI Registry, or inventory, maintained by VITA. According to the AI Utilization Policy, the purpose of the registry and approval process is “to ensure AI is used in a trusted, safe, and secure manner.” VITA identifies some exceptions to its AI registry requirements and approval process, although some of these exceptions are unclear.

VITA’s Enterprise Architecture Standard for AI (EA-225) provides more detailed operational requirements for existing and new uses of AI systems. For example, the standard includes requirements for the approval of new systems, training and awareness programs, model trustworthiness certifications, and controls to mitigate bias risk. According to VITA, the EA-225 standards are intended to provide “requirements for Commonwealth agencies and suppliers on the acceptable and ethical use of AI.” (See chapters 2 and 3 for additional discussion on several problems identified with VITA’s requirements and opportunities to improve the state’s governance of AI use at state agencies.)

VITA’s Office of Data Governance and Analytics (ODGA), established in statute, is intended to “increase the use of data and data analytics to improve the efficiency and efficacy of government services and stakeholder outcomes,” among other purposes. As of August 2026, ODGA offered several services and guidance documents intended to help agencies manage their data and deploy AI responsibly, including data quality scanning, sensitive data scanning, and AI governance and vendor assessment documents.

As of August 2026, there is no comprehensive data on how much VITA-supported agencies or other agencies collectively spend on AI tools, and the primary general-purpose AI tool available to VITA-supported agencies (Copilot Chat) is included as part of the broader Microsoft license.

Mitigating AI-related harms to the public can involve various state entities, including the General Assembly and regulatory agencies

In recent years, states have been considering and passing legislation to balance protecting the public from potential AI-related harms while encouraging adoption of beneficial uses. States' approaches to mitigating potential harms have varied, from broadly applicable transparency and disclosure requirements aimed at AI companies to narrowly tailored laws covering synthetic content.

In Virginia, two legislative agencies, the Division of Legislative Services (DLS) and the Joint Commission on Technology and Science (JCOTS), have roles supporting the General Assembly in evaluating AI technologies and crafting effective legislation to address them. DLS is statutorily required to assist legislators with drafting bills, analyzing the constitutionality or probable legal effect of proposed legislation, and conducting research on proposed legislation. JCOTS is statutorily responsible for advising the General Assembly on technology- and science-related matters and has issued several reports on AI-related topics in recent years. (See Chapter 5 for opportunities for DLS and JCOTS to strengthen support to the General Assembly on AI-related legislation.)

Beyond central IT governance, various executive branch and independent agencies also have roles in addressing AI-related harms under existing statutory provisions that are technology neutral. For example, the State Corporation Commission's Bureau of Insurance regulates unfair discrimination in insurer decisions, regardless of the specific technology used to reach those decisions. In 2024, the Bureau advised insurers that decisions, conduct, or actions impacting consumers that are made or supported by AI must comply with all applicable insurance laws and regulations. Similarly, the Office of the Attorney General has authority to investigate and bring enforcement actions against consumer deception under the Virginia Consumer Protection Act, including instances in which AI may have been used to facilitate the deception.

In September 2026, the governor issued Executive Order 22, establishing an AI task force co-led by the chief transformation officer, counsel to the governor, and secretary of administration. The executive order directs the task force to take several actions, including the following:

- “address the risks posed by AI through executive and legislative action—including risk of workforce displacement and data privacy related to Virginians’ personal information—and develop the expertise, staffing structure, and government systems to most effectively meet the AI moment;”
- “evaluate Virginia law and precedent to identify existing enforcement and regulatory options to address the risks posed by AI growth and the harms being inflicted on all Virginians as a result;”

- “identify and assess the cybersecurity risks posed to Virginia, its state systems, and its critical infrastructure from AI agents, bots, and other developments;” and
- “support Commonwealth-wide efforts to become a leader in secure and responsible use of AI, leveraging recognized best practices and technical experience in governance, privacy, and data security.”

Executive Order 22 also directed the implementation of standards and initiatives related to data centers.

2 Function-Specific AI Use by State Agencies

Function-specific AI tools are designed to perform specific agency operational or administrative tasks, rather than being adaptable across a broad range of uses like general-purpose AI tools (e.g., ChatGPT, Claude, Copilot) (sidebar). Examples of function-specific AI tools used in other states and the federal government include constituent-facing chatbots, fraud detection tools, tools that digitize and classify documents, tools that identify and redact personally identifiable information, and traffic incident detection systems.

In practice, the line between function-specific and general-purpose tools is not always clear, such as with highly capable coding assistance tools.

Function-specific AI may include tools developed by agency staff, custom-built AI solutions specifically developed for an agency by a third party, customized or fine-tuned third-party AI tools adapted for an agency's specific use, commercial off-the-shelf AI tools, or AI functionality embedded in commercial software.

The state's AI Utilization Policy (discussed later in this chapter) requires that executive branch agencies use AI only when it produces a "clear, positive outcome for Virginia's citizens" and is the optimal solution for achieving the intended purpose. Appropriate use cases could include reducing barriers to government services, improving the quality and delivery of services, and reducing bureaucratic delays.

While function-specific AI tools can offer opportunities to improve agency services or processes, they also present potential risks. Examples of potential risks include harmful impacts to citizen rights, benefits, critical services, or safety; exposure of sensitive data; and inefficient use of state funds through unforeseen costs or poor return on investment.

The Virginia Information Technologies Agency (VITA) plays a central role in the governance and risk management of AI for executive branch agencies, including function-specific tools, by approving use cases through its AI registry and by articulating standards and policies for the development, procurement, and use of AI systems.

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

Some Virginia state agencies have deployed or are piloting function-specific AI tools, although most have not done so as of June 2026. However, many of these agencies reported in JLARC's survey of senior IT leaders that they are actively considering or pursuing a function-specific AI tool (sidebar). Given this interest and expanding AI capabilities, both the number of agencies using function-specific AI tools and the variety of use cases will likely grow in the next few years.

JLARC conducted a survey of senior IT leaders at executive branch and independent state agencies. JLARC received completed survey responses from all 80 agencies surveyed (100 percent response rate). See Appendix B for more information.


For this chapter, statistics referencing “agencies” refer to surveyed executive branch and independent state agencies, unless a narrower subset is specified. It does not include higher education institutions, judicial agencies, or legislative agencies.

One-third of state agencies report using function-specific AI tools

As of June 2026, a minority of state agencies were using an AI tool for a specific agency function. About a third of agencies reported using or piloting at least one function-specific AI tool (Figure 2-1) (sidebar). However, an additional 30 percent of agencies were actively considering or pursuing adoption of at least one tool.

FIGURE 2-1

One-third of agencies are currently using or piloting at least one function-specific AI tool



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

NOTE: “Not currently using or pursuing a function-specific AI tool” includes agencies that indicated they were not using, piloting, or actively considering or pursuing the adoption of at least one function-specific AI system or tool.

The remaining 38 percent of agencies reported they were not currently using or considering a function-specific AI tool. These agencies most frequently cited not having identified a business need or clear use case for AI (72 percent). Other top reasons were concern about data privacy risks associated with AI (31 percent) and insufficient technical skills among staff to build, deploy, or maintain AI systems (17 percent).

Agencies using function-specific AI are both procuring these tools from outside sources and developing them internally. Eighty-eight percent of agencies using function-specific AI reported using or piloting at least one tool procured from a third-party vendor, and 65 percent reported using or piloting at least one tool developed entirely or partially by agency staff.

Larger agencies were more likely to have adopted a function-specific AI tool, likely reflecting their greater technical and financial resources. Eighty-five percent of agencies using these tools had more than 200 full-time equivalent salary and wage employees, and 46 percent had more than 500.

Function-specific AI use cases varied across agencies (sidebar). Examples of tools included:

- a license plate text review tool that flags potentially offensive content before approval of personalized license plate requests (Department of Motor Vehicles, fully deployed);
- a tool that helps job seekers tailor resumes and cover letters to specific job descriptions (Virginia Works, pilot);
- a tool to analyze public comments received during the Petition for Rule-making process (Virginia Department of Health, fully deployed);
- a tool that identifies potential security threats in scanned body and bag images (Department of Corrections, fully deployed);
- a tool that analyzes data across siloed systems to improve project cost estimation (Virginia Department of Transportation, pilot);
- a financial management tool that integrates and streamlines processes such as purchase requisitions, budget modifications, invoice processing, and financial reconciliation (Virginia Department of Emergency Management, fully deployed);
- an agentic permit application review tool (Department of Environmental Quality, pilot);
- an agentic tool to help build data documentation (i.e., data dictionaries) (Department of Behavioral Health and Developmental Services, pilot); and
- an AI capability embedded in a recruitment management system that generates standardized resume summaries of applicants relative to job descriptions (Department of Human Resource Management, fully deployed).

JLARC staff reviewed how executive branch and independent agencies and authorities are using AI tools, as well as barriers they face. However, **evaluating whether AI use cases were effective or appropriate was beyond the scope of this study.**

Majority of agencies using or considering function-specific AI tools aim to improve agency efficiency, data quality, or public services


Agencies reported using or considering function-specific AI to support both administrative operations and programmatic work. Agencies most commonly reported using or considering function-specific AI tools to support information technology (56 percent), mission-specific program or service areas (52 percent), and finance and accounting (30 percent) (Figure 2-2) (sidebar).

Agencies reported a variety of reasons for using function-specific AI. Of the agencies currently using or pursuing at least one function-specific AI tool, the most frequently cited anticipated benefits were more efficient internal operations, cost savings, or staff productivity (84 percent); improved data quality or accessibility for analysis and decision-making (56 percent); and improved public services, such as constituent-facing services and interactions (54 percent).

Potential examples of AI use cases in information technology include improving cybersecurity, monitoring system performance, and assisting with software development.

FIGURE 2-2

Agencies most frequently report using or considering function-specific AI to improve internal efficiency and support agency IT



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

NOTE: Fifty of 80 surveyed agencies reported using or actively considering at least one function-specific AI tool. Agencies could select more than one anticipated benefit and agency function, so percentages total more than 100.

State's AI requirements lack risk-based criteria, clarity, and timely updates

VITA governs AI use by executive branch state agencies primarily through its AI Utilization Policy and the Enterprise Architecture Standard for AI. The AI Utilization Policy focuses on broad principles for agency adoption and use of AI, while the Enterprise Architecture Standard generally addresses technical and operational requirements for AI systems.

These requirements apply to all AI systems except for a few noted exceptions (e.g., cybersecurity systems). They are not exclusive to function-specific AI, although the requirements are primarily oriented around the development, procurement, or use of function-specific AI systems.

Other technology-neutral state requirements also apply to AI systems and the data that interacts with those systems, such as VITA's Information Security (SEC530) and Data Classification (SEC540) standards.

As the state continues to navigate the rapidly evolving risks and opportunities associated with AI, national and international guidelines (e.g., National Institute of Standards and Technology [NIST], International Organization for Standardization, International Electrotechnical Commission) indicate policies governing AI use should have several key qualities:

- Risk-tiered: standards establish meaningful safeguards for higher-risk uses to protect the public and the state from harm without imposing overly burdensome requirements on beneficial lower-risk applications.
- Clear and actionable: requirements are specific and practical, giving agencies clear direction on expectations and how to achieve compliance.
- Up to date: standards keep pace with evolving AI capabilities, use cases, and risks.

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

AI risks vary widely in type and severity depending on their purpose, the context in which they are used, data sensitivity, and access to organizational files and systems. For example, an AI tool that flags potential data entry errors for human review poses less risk than one that autonomously manages traffic patterns.

National and international guidelines (e.g., NIST, Organisation for Economic Co-operation and Development [OECD], U.S. Office of Management and Budget) recommend that organizations prioritize AI risk management policies and resources based on each AI tool's risk likelihood and severity. NIST notes that attempting to eliminate all AI risks is impractical, inefficient, and can waste limited resources.

At least eight other states (e.g., California, Indiana, Colorado, Idaho, Washington, Maryland, Oregon, Texas) differentiate some of their AI requirements based on risk. While their risk tiering approaches vary, these states categorize AI uses or systems into risk levels (e.g., low, medium, high) and apply differentiated requirements for approval, use, or oversight based on the assessed risk level. (See Appendix F for examples of other states' risk-based approaches to AI oversight.)

While Virginia's AI standards contain key risk mitigation measures that align with widely cited guidance (e.g., NIST, OECD), most requirements apply uniformly regardless of risk level. For example, almost all AI systems and tools require:

- approval from senior IT agency staff, certain VITA divisions, the chief information officer (CIO) of the Commonwealth, and the secretary of administration as part of the AI registry process (sidebar);
- annual certification for trustworthiness according to NIST-defined characteristics (e.g., valid, reliable, safe, secure, resilient, accountable, transparent, privacy-enhanced, explainable, interpretable, and fair);

VITA staff report denying only one proposed use case submitted to the AI registry. However, VITA indicated that they may work with some agencies to refine their planned use cases rather than deny them.

- internal and external validation of system design, data collection, and measurements in a system’s deployment context; and
- ongoing evaluation that includes testing for degradation in model performance, risk assessment, incident tracking, and evaluating user experience and process integration.

“*[The state should] consider establishing a documented exception or expedited review process for existing commercial solutions with embedded AI capabilities or low-risk implementations. This would provide agencies greater flexibility while maintaining appropriate governance.*”

– Agency senior IT leader

Largely uniform requirements for AI tools can delay or discourage the adoption of lower-risk AI tools and create unnecessary burdens on VITA and agency staff. For example, using AI-powered optical character recognition to translate handwritten documents into digital text that can be searched and edited is typically a low-risk use case that should not be burdensome to implement.

Uniform requirements also may not impose sufficiently stringent safeguards on higher-risk use cases, and some states have established more rigorous requirements for such uses. For example, Maryland imposes more extensive requirements for the deployment of high-risk systems, including a comprehensive algorithmic impact assessment and the identification of robust safeguards that sufficiently address the relevant risks. Similarly, Washington requires agencies implementing high-risk AI to develop an adoption plan that includes risk mitigation strategies, data quality standards, metrics for measuring performance, and a formal process to regularly monitor outputs after deployment.

In addition to having uniform requirements, VITA has not provided agencies with guidelines for assessing AI risk. Although the AI Utilization Policy requires agencies to identify AI risks as part of their internal approval prior to submitting use cases to the AI registry, it does not specify criteria for doing so. VITA also recently amended this policy to require agencies to annually audit high-risk systems for compliance with several data-related requirements but does not define “high-risk.”

VITA should develop a standardized AI risk assessment instrument for agencies to evaluate potential use cases. This instrument could build on the data sensitivity and critical business function criteria used to designate IT systems as sensitive in VITA’s Information Security Standard, but it should also incorporate AI-specific considerations such as the potential impact of decisions made or informed by AI. At minimum, the instrument should:

- provide clear criteria that assess risks related to data and system security, as well as policy-related risks (e.g., access to critical government benefits and services; impacts to citizens’ rights, health, or safety);
- use a standardized assessment methodology, such as questions with yes/no or categorical answers rather than open-ended questions (e.g., Does the proposed AI use involve confidential or personally identifiable information? Is the proposed AI use public-facing?); and

- produce a risk designation for a potential use case (e.g., low, medium, high, unacceptable risk) based on the agency's responses to these criteria, while allowing agencies to assign a higher risk level based on their specialized knowledge of the policy area, population served, or business process.

Other states' approaches provide insights into how to classify AI risk. For example, common criteria for identifying higher risk systems and use cases include whether they involve confidential or personally identifiable information; disproportionately impact legally protected or vulnerable populations; directly or indirectly impact public rights, safety, or benefit eligibility; affect critical government operations and services; or influence other consequential decisions. Common characteristics associated with lower risk systems or use cases include using public or non-sensitive data, performing purely internal tasks, and presenting minimal risk of harm to individuals or agency operations.

VITA should consult with subject-matter experts, such as higher education institutions with relevant expertise in AI policy, to ensure the risk assessment instrument contains the appropriate criteria and can be easily understood and applied by agencies.

VITA should also revise its AI standards to reflect varying approval, use, and monitoring requirements needed to mitigate different levels of risk, as identified using the risk assessment instrument. For lower risk uses, the requirements should place fewer burdens on agencies and necessitate minimal oversight from VITA. As part of this effort, VITA should revise the AI use case approval process so that agencies can approve low-risk AI use cases, while still being required to report those use cases on the AI registry so their use is known. Many other states have a centralized AI inventory, but some do not require centralized use case approval (e.g., New York) or do not require centralized review for low-risk uses (e.g., California). VITA could retain the right to audit any registered use case if it disagrees with the agency's risk assessment.

In contrast, higher-risk uses should be subject to greater requirements and attention from VITA and, as appropriate, other state entities. For example, use cases that are high-risk because of their potential policy implications (e.g., those impacting citizens' rights, benefits, or safety) may need to be escalated for specialized higher review and approval, such as by the relevant agency secretariat or the Office of the Attorney General. Beyond the initial risk assessment, agencies seeking to implement high-risk AI uses should be required to demonstrate they have thoroughly assessed the potential harms of the proposed use case and taken steps to adequately mitigate those risks. Higher-risk use cases should also be subject to enhanced monitoring and oversight to ensure the AI tool is performing as intended and unintended harms are identified.

VITA should also clearly define its oversight role in ensuring agencies remain in compliance with AI requirements. Compliance mechanisms should be proportionate to risk, with greater attention on high-risk uses.

RECOMMENDATION 1

The Virginia Information Technologies Agency should develop and publish an artificial intelligence (AI) risk assessment instrument that provides agencies with clear, standardized criteria for evaluating potential AI use cases and determining their risk level.

RECOMMENDATION 2

The Virginia Information Technologies Agency (VITA) should develop and implement a risk-tiered approach for governing executive branch agencies' use of artificial intelligence (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 Artificial Intelligence Utilization Policy and the Enterprise Architecture Standard for Artificial Intelligence to reflect tiered requirements.

Unclear AI standards make compliance challenging for agencies

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 governing state agency AI use as a major or moderate constraint on their pursuit or implementation of AI. For example, while the standards require registration and approval of AI use cases, only a minority of executive branch agencies found these requirements to be “very clear,” including which use cases must be submitted to the registry (30 percent), the specific information about each use case that must be submitted (32 percent), and VITA’s approval criteria for use cases (21 percent).

JLARC staff’s review of the AI standards supports agency concerns. Multiple provisions of the standards are unclear about requirements and how agencies can meet them. For example:

“
Clearer guidelines would
resolve major
operational uncertainty.”
”

– Agency senior IT
leader

“
The compliance
requirements are both
vague and vacillating –
amorphous at best and
obstructive at worst.”
”

– Agency senior IT
leader

- “COV AI Systems shall be subject to Test, Evaluation, Verification, and Validation (TEVV) review throughout the AI System lifecycle,” but VITA staff indicate they have been unable to identify resources that could provide a defensible TEVV approach for agencies.
- “COV AI Systems shall be annually certified for trustworthiness according to... characteristics defined by NIST,” but VITA identifies no certifying body or acceptable methodology.
- “COV AI Systems shall employ data quality tools that can score the data and shall meet a minimum accuracy threshold of 90%,” but VITA does not specify what data should be evaluated (e.g., training, operational input, model output), identify acceptable tools or methodologies for determining accuracy, or specify how often data should be evaluated.

- “COV AI System models shall be capable of providing explanations or insights into their decision-making processes to facilitate trust and understanding;” however, this is technically infeasible for many modern AI systems because explanations may be unreliable, omit key information, or not reflect the actual process that produced the output.
- Multiple agencies have expressed confusion about what qualifies as “AI embedded within common commercial products,” to which the AI standards do not apply.
- While VITA has told agencies that registry approvals apply to individual use cases rather than blanket approvals of AI tools or systems, the AI standards do not make this distinction clear (e.g., “[Agencies] must complete the following approval process before deploying any AI system”).

Unclear requirements may lead agencies to interpret compliance differently, undermining state efforts to ensure AI systems are safe and effective. For example, because VITA’s AI standards do not clearly state that *each individual use case* must be registered rather than *AI systems*, agencies may use previously approved systems for higher-risk uses without additional oversight. Because some AI tools can be configured for a variety of purposes, the same tool can be applied to both low- and high-risk use cases (e.g., analyzing and summarizing population-level crime statistics versus summarizing an inmate’s criminal history for parole determinations). Agencies may also disregard some of VITA’s requirements if they cannot determine how to comply.

VITA should revise the Enterprise Architecture Standard for AI and the AI Utilization Policy to ensure they are as clear and actionable as possible. Some states (e.g., Maryland, Oregon) have created implementation guidance with their AI requirements to help agencies apply them. VITA should consult with agencies to identify where additional guidance is needed and publish this guidance on its website.

RECOMMENDATION 3

The Virginia Information Technologies Agency should (i) revise the Enterprise Architecture Standard for Artificial Intelligence and the Artificial Intelligence 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.

State’s AI standards have not kept pace with emerging risks

AI capabilities are advancing rapidly, and the state’s AI standards need to keep pace with technological change and evolving risks. In particular, agentic AI capabilities have improved dramatically in the last two years (sidebar). For example, top agents’ performance on a benchmark evaluating open-ended task completion improved from 12 percent in 2024 to 66 percent in 2025, nearing human baseline performance of 72 percent. Agentic capabilities are increasingly being incorporated into commercial

“The standard is relatively high-level and principle-driven, which means it reads well but leaves a lot of execution detail undefined...By adding more complexity (new ambiguous standards) and less actual tools/resources, VITA is directly and indirectly negatively impacting agency ability to meet taxpayer expectations and obligations, rather than enabling the mission of agencies.”

– Agency senior IT leader

AI agents are AI tools that can autonomously plan and take actions to achieve a specified goal without needing a human to define each step to reach that goal. If granted access, agents can take actions such as modifying or moving files, creating new documents or databases, sending emails, searching the internet, or performing specialized tasks.

Because agents may take actions unintended by the user, they can pose greater risks than non-agentic AI systems.

products, making powerful tools available to a broader population of users. Some Virginia state agencies are exploring use cases involving agents and have received approval from VITA through the AI registry.

VITA's AI standards have not sufficiently kept up with emerging risks, especially those related to agentic AI. The standards do not refer to agents or specify how safeguards extend to autonomous tools. While some existing requirements may sufficiently apply to agents, other key safeguards should specifically address the increased risks they introduce. For example, because agents can independently determine how to carry out instructions, the standards should address appropriate limits to agents' autonomy, including when human oversight is required.

Without standards that adequately address agentic risks, agencies may deploy these tools without necessary safeguards against unintended agent behavior. Agents pose heightened risk because they can independently take actions in multi-step workflows. Early errors can compound across subsequent steps; agents may act beyond their intended scope; and overly broad permissions or multiple interacting agents can further amplify these risks.

VITA's most recent Information Security Report (2024) identified access controls as the most common IT audit and risk finding for executive branch and independent agencies, and the Auditor of Public Accounts identified 17 findings related to access controls across examined agencies in its most recent Single Audit Report for the Commonwealth.

“The registry process and standards just seem like a paperwork exercise... they haven't updated the policy in three years. And it's a black hole. No one in the process is checking back about how we intend to ensure compliance.”

– Agency senior IT leader

VITA's AI standards also do not sufficiently address risks associated with agents' or other AI tools' access to state systems and data. AI tools can inherit users' access permissions or be granted broad access during configuration, making it easier for sensitive information to be retrieved by unintended users. An agent operating with those permissions can also modify, delete, or share files. Recent state reports highlight ongoing challenges with overly broad access controls across agencies (sidebar), and VITA staff cited a lack of adequate access controls on agency files as a key reason for limiting the features of the enterprise general-purpose AI tool. Various strategies could be considered to help safeguard agency systems and data, including requiring that AI tools be configured by default with the most restrictive access settings that still allow intended functionality; granting agents temporary credentials to sensitive systems and data that expire after task completion; requiring certain data to be accessed only through centrally managed connection points, such as Model Context Protocol servers; and requiring agencies to use sensitive data locator tools—available through VITA's Office of Data Governance and Analytics—before approving certain high-risk use cases.

These governance gaps exist in part because VITA has not updated its AI standards frequently enough to reflect AI's rapid technological advancements. The Enterprise Architecture Standard for AI was last updated in November 2023. The AI Utilization Policy was recently updated in June 2026, but this revision was released over two years after the policy was first created, and it included only a few substantive changes. VITA staff indicate the slow pace of updates is due to competing priorities and its lengthy internal review process.

VITA staff acknowledged that the existing AI standards do not adequately address all emerging risks, in particular agentic AI. VITA should prioritize updating the Enterprise Architecture Standard for AI, and publishing additional guidance as needed, to

ensure the state has sufficient safeguards to mitigate risks related to agentic AI and any other emerging risks identified. Staff indicated that the standard is currently being revised and will include requirements specific to agents. Any revisions to this standard will be available for public comment before adoption.

Going forward, as AI capabilities continue to evolve and state agencies continue to adopt new AI use cases, VITA should review all AI-related standards, policies, and guidance provided to agencies at least annually to ensure they remain up-to-date and adequately address emerging risks and any gaps in state governance.

RECOMMENDATION 4

The Virginia Information Technologies Agency should (i) conduct a one-time targeted review to update the Enterprise Architecture Standard for Artificial Intelligence 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.

VITA's review and approval processes beyond the AI registry have a reasonable basis but can be burdensome and are not well coordinated

As referenced previously, approval of an AI use case begins with the AI registry process. An agency submits a request through the registry with general information including the purpose, sensitivity of the system's data, and model details. Staff at two VITA divisions and the CIO review the use case submission and determine whether to approve it. VITA policy also requires an additional level of review and approval by the secretary of administration.

The AI registry, however, is only the first of several required review and approval processes for new AI projects or procurements, which can add substantial time and work to implement an AI solution. The number and intensity of reviews required for proposed AI projects vary depending on the cost, procurement type, and risk level of the procurement (sidebar).

Documenting the full sequence of project approval and procurement processes for IT projects, including AI projects, is challenging, as not even VITA has a comprehensive process map. However, the general review and approval processes are as follows:

- **Business case request:** The agency must first request and justify the project or procurement through a formal investment business case. This step alone is subject to seven levels of review prior to approval.
- **Business case approval:** Once the business case is developed, the agency must submit a formal approval action. Low-risk projects go through two levels of review; higher-risk projects can require three levels of review.

Risk in IT procurements is based on the cost and terms of the contract (e.g., involves two or more state bodies, initial contract term exceeds five years). This is a separate consideration from risk related to the adoption of an AI tool.

- Project governance request: If a project will cost more than \$250,000, the agency must also submit a project governance request, which informs the level of oversight required by VITA and is reviewed by a single division.
- Solicitation or contract review: Once these steps are complete, the agency can move forward with procurement. Solicitations or contracts can be subject to as many as seven additional levels of review if valued over \$1 million.
- Cloud service review (COV Ramp): For any cloud-based service (most AI tools), an additional review is required, consisting of three levels.

Review timeframes vary widely, with individual steps that may last multiple months. For example, the initial phase of the COV Ramp review alone averaged over 50 days in 2025. Taken together, a high-risk IT procurement project could require 20 reviews during the project and procurement approval cycles, including multiple reviews by the same division. For instance, the enterprise architecture and cybersecurity and risk management divisions each review a procurement or project four separate times in this process.

Several indications suggest these reviews are not well coordinated and may be burdensome for agencies trying to adopt AI tools. In interviews and survey responses, agencies' senior IT leadership reported frustration with the project review and approval processes and expect that AI will only escalate inefficiencies.

“
VITA traditionally has a
hard time talking to
VITA.
”
— VITA senior staff

“...in my past experiences with procurements, I find the state and VITA processes to be very difficult and unclear. And too often the VITA procurement, security, contracts, and state legal, are not on the same page and don't know what the others are doing. I'm not sure how I am supposed to know things when they themselves don't either.”

“[We need] a better understanding of the approval process, and clear guidance on prioritization of systems requirements. As it stands, it seems like VITA leadership doesn't understand their own procurement processes for AI.”

“A lot of these issues and constraints are not specific to AI—they apply to non-AI technology products, procurements and projects as well. The rapid rise and escalation of AI and AI-enabled IT solutions, has however, further exacerbated endemic concerns, pain points and issues experienced by agencies.”

VITA leadership is aware of issues with multiple separate approvals needed for new AI projects or procurements and indicate they are taking steps to reduce the number of approvals needed to implement some AI use cases. Most notably, VITA is working to develop statewide contracts for major AI companies, such as Anthropic and OpenAI, which reportedly will allow agencies to develop and deploy AI use cases using these products without going through IT procurement-related reviews. However, any AI use case involving a third-party AI product without a statewide contract will still be subject to the review processes outlined previously.

Fully integrating state project and procurement processes is beyond the scope of this project and involves various legal and technical considerations. While statute requires certain project review steps, VITA has discretion on the structure of these reviews. VITA leadership should document the complete, step-by-step AI project and procurement processes to provide internal clarity and ensure state agencies can easily navigate all applicable requirements. In the near future, VITA should consider having a third party evaluate the processes to ensure they are as efficient as possible while still appropriately mitigating risks.

RECOMMENDATION 5

The Virginia Information Technologies Agency should develop and publish a comprehensive process map detailing the required review and approval steps for state agency artificial intelligence projects and procurements.

Vendors are adding AI capabilities into existing state software, raising risk of improper data use

A common concern raised by agencies in interviews and survey responses is the introduction of AI capabilities into non-AI commercial software already in use, which can introduce AI-related risks into these systems. Forty-four percent of agencies indicated that at least one AI feature had been added to a non-AI software or system already used by the agency.

AI functionality added to software without authorization introduces unmanaged risk. Staff may begin using these features (in some cases unknowingly) before the use case is approved or the functionality is vetted for compliance with the state's AI standards. In addition, data interacting with the software's AI feature may be used for model training, sent to external servers or third-party model providers, or exposed in other ways.

VITA has taken steps to mitigate risks related to vendor-embedded AI by developing standard contractual protections for all cloud-based software. These terms will be used in new cloud contracts and recompetes of existing contracts, requiring that: (1) any AI functionality can be disabled by VITA or the agency; (2) vendors must provide at least 30 days' notice before adding new AI features, which remain disabled by default unless authorized in writing by VITA or the agency; (3) vendors and their sub-processors cannot use state data to train AI models; and (4) vendors must publish and maintain up-to-date model system cards (i.e., technical and performance documentation) for any AI system used.

VITA should follow up with agencies within a year to ensure this language is sufficiently preventing and mitigating risks related to the introduction of AI capabilities in systems where these contractual terms have been applied.

3 General-Purpose AI Use by State Employees

General-purpose AI tools respond to user prompts or instructions by generating content, such as text, images, and code, or taking actions (e.g., scheduling a meeting and sending invitations), and can be used across a broad range of tasks and contexts. These general-purpose AI tools include chatbots, such as Copilot, Claude, and ChatGPT, as well as tools tailored to specific domains or output types that can still support a broad range of tasks. For example, coding assistants (e.g., Claude Code, OpenAI Codex, GitHub Copilot) can perform tasks such as writing, testing, and troubleshooting code; producing documentation for existing codebases; and migrating data, code, or applications to new systems.

When used responsibly, general-purpose tools may improve employee productivity and work quality in certain applications. By assisting with completing mundane tasks more quickly, such as drafting and editing routine documents, employees have more time to spend on complex, creative, specialized, or face-to-face work.

If not used responsibly, general-purpose AI may increase various risks, such as undetected and consequential errors in government work. Errors can include fabrications presented as facts, flawed or biased analyses, or false citations that go unnoticed without careful human review. Available AI benchmarks indicate even leading models have a high rate of false or flawed outputs in certain situations, such as answering open-ended knowledge questions (sidebar).

Without appropriate safeguards, general-purpose AI can also risk exposure of sensitive data. Mitigating this risk is especially important given the volume of personal and legally protected information held by the government.

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

Almost all executive branch and independent state agencies (97 percent of agencies surveyed by JLARC) allow staff to use general-purpose AI tools, although many have adopted restrictions (sidebar). Sixty-two percent of surveyed agencies allow staff to use only agency-approved general-purpose AI tools, 27 percent limit staff use to approved use cases, and 11 percent allow only authorized staff to use them.

The majority of state employees at executive branch agencies served by VITA appear to use general-purpose AI tools, although usage data is limited. Between July 27, 2026, and August 27, 2026, available data indicates an estimated 76 percent of employees (41,556 out of ~55,000) used a general-purpose AI tool to some extent.

General-purpose AI tools may generate inaccurate outputs that are presented as factual. For example, a benchmark cited in Stanford Institute for Human-Centered AI's 2026 AI Index Report evaluated the accuracy of 26 models' responses to 6,000 knowledge questions. The rate of false answers (hallucination rate) ranged from 22 to 94 percent across models.

JLARC conducted a survey of senior IT leaders at executive branch and independent state agencies. JLARC received completed survey responses from all 80 agencies surveyed (100 percent response rate). See Appendix B for more information.

VITA and agencies report that staff use general-purpose AI tools for a variety of purposes, including improving code, summarizing documents or notes, and drafting emails. Comprehensive data on how state employees are using these tools does not exist.

State employees' undirected use of publicly available general-purpose AI tools can increase risk of data exposure


VITA is taking steps to make additional general-purpose enterprise tools available to state employees. VITA is currently negotiating contractual terms with Anthropic, OpenAI, and Google to potentially establish statewide contracts agencies could use to purchase licenses for Claude, ChatGPT, and Gemini.

VITA has approved a small number of agencies to use general-purpose AI tools other than Copilot, typically for limited purposes. For example, the Department of Motor Vehicles received approval to use ChatGPT to flag potentially offensive content in personalized license plate requests.

In August 2025, VITA started providing an enterprise version of Copilot to all the agencies it supports. While Copilot is currently the only enterprise general-purpose AI tool available to all executive branch employees (sidebar), state employees can also access publicly available tools, since many are freely available online.

AI usage data indicates state employees frequently use general-purpose AI tools other than the VITA-provided enterprise tool. While many state employees at VITA-supported agencies used the enterprise Copilot to some extent from July 27, 2026 to August 27, 2026, it accounted for only a small proportion of the total requests and data exchanged with general-purpose AI tools during this period, according to available data from VITA (Figure 3-1). For example, during this period, employees made more requests to and exchanged more data with OpenAI's ChatGPT than the VITA-provided enterprise Copilot (sidebar).

FIGURE 3-1
State employees at VITA-supported agencies interact frequently with publicly available AI tools



SOURCE: State employee usage of AI services data for July 27, 2026, to August 27, 2026, provided by VITA.

NOTE: "Total data exchanged" includes all data between the user and server in both directions (e.g., uploads, downloads, prompts, and responses). "Other general-purpose AI tools" aggregates usage data for ChatGPT (OpenAI), Claude (Anthropic), Gemini (Google), and the consumer version of Microsoft Copilot. Usage data for general-purpose AI tools with fewer than 1,000 users is not included.

State employees' reasons for using publicly available tools likely vary but may be driven partially by limitations in the available enterprise tool and VITA's unclear state policy guidance. Some agencies report being dissatisfied with limitations of the VITA-approved version of Copilot (e.g., limited ability to upload files, lack of integration with

other Microsoft applications). In addition, multiple VITA staff reported hearing that at least some state employees have been using publicly available tools for their coding capabilities. Additionally, as discussed later in this chapter, VITA has not clearly communicated its overall policy on whether state employees can use publicly available, general-purpose AI tools.

Although specific protections may vary, employer-provided enterprise AI tools generally are more secure than publicly available tools because they can provide increased data protections (e.g., data is not used to train the AI model), greater administrative oversight, and additional configuration options (e.g., ability to disable features). In contrast, data entered into publicly available tools may be used to train future models, retained for extended periods, transmitted outside of the state's secure environment, or shared with other third parties, increasing the risk of data exposure or breach if not used responsibly.

Available data does not indicate how Virginia state employees are using publicly available general-purpose AI tools, and not all uses of these tools carry the same level and type of risk. Some states (e.g., North Carolina, Minnesota, Vermont, New Jersey) allow employees to use publicly available general-purpose AI tools within defined parameters, such as restricting the data that can be entered and requiring all accounts to be registered with a state email address for records retention.

Currently, however, Virginia is operating at a higher-than-necessary risk of state employee misuse of general-purpose AI tools. As noted in the next section, the state's risk is elevated because of several factors, including the lack of a statewide acceptable use policy and staff training on responsible AI use, as well as a lack of clarity on which specific tools may be used for which purposes. Recommendations 6, 7, and 8, if implemented, would help address these gaps and reduce the state's risk associated with general-purpose AI tools.

Policy and training gaps for general-purpose AI tools increase state's risk of sensitive data exposure and other problems

As governments and private organizations adopt general-purpose AI tools, many are also implementing strategies to manage associated risks and promote responsible use. This includes creating policies to establish clear expectations for acceptable staff use of these tools and training to build understanding of their capabilities and limitations.

Agencies surveyed by JLARC generally agreed that implementing these strategies statewide would help support effective and responsible staff use of general-purpose AI tools.

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

AI agents are AI tools that can autonomously plan and act to achieve a goal. Because agents may take actions unintended by the user, they can pose greater risks than non-agentic AI systems. These risks are generally greater when agents have broader access to systems and files. See Chapter 2 for more information on risks associated with AI agents.

Establishing acceptable use policies for general-purpose AI tools can help organizations mitigate risks associated with these tools. Without clear requirements, employees may be more likely to unintentionally share sensitive information with AI tools that lack proper safeguards, rely on AI-generated content without adequate oversight, grant agents excessive permissions to access and act on files and systems, or use tools in ways that conflict with legal or regulatory requirements (sidebar). In addition, addressing employee misuse may be more difficult if responsibilities for using these tools are not clearly defined.

Other states (e.g., Oregon, North Carolina, Connecticut, California, Vermont, Hawaii, South Dakota, Minnesota, Maryland, Idaho) have developed acceptable use policies or guidance outlining expectations for state employee use of general-purpose AI tools. These policies include requirements related to data entry, employee responsibility for reviewing and validating AI-generated outputs used in work products, incident reporting, records retention, training prior to use, and prohibited uses.

Virginia has little formal statewide guidance for state employees on the use of general-purpose AI. While a few provisions in VITA's existing AI standards apply to general-purpose AI use, such as limiting sensitive data entry to private AI systems solely accessible by Commonwealth users, the standards are primarily requirements for agencies related to the development, procurement, and use of function-specific AI tools.

VITA has also not clearly communicated its policy on state employee use of publicly available general-purpose AI tools or uses of enterprise tools that still require registry approval. VITA staff have indicated to JLARC that only enterprise tools can be used for work-related queries and that any use involving sensitive data or integration into an agency business or operational process must be registered, even when using the VITA-provided enterprise tool. However, these restrictions are not clearly stated in VITA's AI standards, and agencies have expressed confusion in survey responses and VITA-led forums about which general-purpose AI tools can be used by staff and which use cases require approval through the registry process.

Most agencies have also not yet established their own policies on general-purpose AI use for their staff, and agency IT leaders indicated a desire for statewide guidance. As of June 2026, only 28 percent of surveyed agencies reported providing written policy or guidance for staff AI use. Most surveyed agencies (90 percent) indicated that a statewide acceptable use policy for general-purpose AI tools, with the ability for agencies to establish additional restrictions, would be very helpful (70 percent) or somewhat helpful (20 percent) in supporting effective and responsible use by staff (Figure 3-2).

FIGURE 3-2

Agencies indicate a statewide acceptable use policy for general-purpose AI tools would help support effective and responsible use by staff



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

VITA should develop an acceptable use policy governing state employees' use of general-purpose AI tools (sidebar). At minimum, the policy should clearly articulate:

- whether, and under what conditions, state employees may use publicly available general-purpose AI tools;
- which use cases must be registered even when using approved or VITA-provisioned tools;
- prohibited use cases;
- requirements related to the use of state data and information;
- employee obligations for use and oversight of AI-generated outputs;
- any restrictions related to granting AI tools access to state files, systems, or applications; and
- whether, and under what conditions, state employees may use agentic capabilities of general-purpose AI tools that can autonomously take actions, such as modifying files or sending emails.

The state already provides guidelines for employee use of other electronic tools. The state has an acceptable use policy for electronic communications, internet use, and social media, which agencies are permitted to supplement with additional requirements or context, and which all state employees must review and acknowledge.

This policy should be reviewed annually, updated as needed, and require employee attestation that they have read and will abide by the policy. Executive branch agencies should be required to either adopt VITA's acceptable use policy or develop an agency-specific policy that meets or exceeds VITA's baseline requirements. Agencies may add further restrictions, as appropriate, to mitigate risks that might be specific to their agency.

RECOMMENDATION 6


The Virginia Information Technologies Agency (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 the 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 during initial onboarding and annually thereafter.

In addition to VITA's lack of clarity on its policy limiting state employees to use of enterprise general-purpose AI tools, VITA does not provide clear guidance on which tools are approved for use by agency staff. VITA directs agencies to its online "AI Roadmap" for approval statuses of general-purpose and other AI tools; however, the roadmap is confusing and could be easily misinterpreted (Figure 3-3). For example:

- The term "AI Roadmap" does not clearly indicate to state employees that it includes a list of approved general-purpose AI tools.
- As of September 2026, ChatGPT is listed as "Approved," but hovering over the entry reveals an "Approved Range" that ended in June 2026, indicating the approval has expired.
- The roadmap and other published information from VITA about AI use do not clearly define "approved." For example, VITA staff provided JLARC with different interpretations of what ChatGPT's approved status means: (1) approved only for enterprise accounts, although there was no statewide contract from which agencies could purchase enterprise ChatGPT licenses during the two years it was approved; and (2) approved only for certain agencies for certain use cases.

FIGURE 3-3

VITA's AI Roadmap lists ChatGPT as approved, but approval expired in June 2026



SOURCE: JLARC screenshot of VITA's AI Roadmap (September 11, 2026).

VITA should publish clear and accessible information on its website about which general-purpose AI tools state employees may use. Any restrictions on how state employees can access approved tools should also be noted (e.g., only through enterprise accounts) and remain consistent with the acceptable use policy. VITA should also ensure this information remains up to date when new general-purpose tools are approved or if a tool's approval status changes.

RECOMMENDATION 7

The Virginia Information Technologies Agency 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.

Lack of training on capabilities, limitations, and risks of AI can increase errors and misuse of these tools

A common AI risk management practice in other states is training staff on the responsible use and limitations of AI tools (sidebar). Without training, some employees may not be sufficiently aware of AI tools' limitations, increasing the risk of accidental misuse. Employees may not verify AI outputs, input sensitive data into public tools, or be unaware of laws and regulations relevant to AI use (e.g., privacy laws, records retention requirements). Other staff may be overly hesitant to use general-purpose AI tools because they are unaware of their capabilities, how to use them effectively, or how to mitigate their risks. Additionally, without training staff on the responsible use and limitations of general-purpose AI tools, agencies may find it more challenging than necessary to hold employees accountable if these tools are misused.

VITA's AI standards already require agencies to provide training and awareness programs for staff involved in "AI-related activities"; however, most state agencies have not yet implemented general-purpose staff AI training. As of June 2026, only 13 percent of surveyed agencies offered basic staff training on responsible general-purpose AI use, and only 4 percent offered role-specific training.

In interviews and survey responses, IT leaders at various state agencies and VITA staff identified insufficient AI literacy among agency staff as a leading constraint to their use of AI. Among agencies currently using or pursuing AI, 76 percent reported that insufficient staff AI literacy was a major (34 percent) or moderate (42 percent) constraint on their agency's ability to pursue or use AI tools or systems over the past 12 months.

VITA staff report they are considering partnering with a nonprofit AI training provider for public service staff that is commonly used by other states and government entities to offer AI training to state employees. Training topics have not been finalized, but VITA is considering both general and role-specific training (e.g., for developers, managers). VITA staff report that this training would be voluntary.

Because general-purpose AI tools are accessible and widely used among state employees, state employees should be required to complete annual AI training to ensure they have basic knowledge of how to use these tools responsibly (sidebar). Most agencies (92 percent) indicated that general training for all state employees would be very helpful (72 percent) or somewhat helpful (20 percent) in supporting effective and responsible general-purpose AI use (Figure 3-4).

Many other states require or offer AI training for state employees.

Some states (e.g., Texas, Washington, Maine) require staff to complete AI training annually. Others require that employees complete training prior to using AI tools (e.g., Oklahoma, Delaware) or certain types of general-purpose tools (e.g., North Carolina, Oregon). Some states (e.g., Idaho, California) also require or offer role-based AI training, such as for senior leadership, other management staff, and technical staff.

Virginia law already requires VITA to develop annual cybersecurity training, in consultation with the heads of executive branch and independent agencies and representatives of the chief justice of the Supreme Court and the Joint Rules Committee of the General Assembly.

FIGURE 3-4**Agencies indicate staff training would help support effective and responsible general-purpose AI use**

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

At a minimum, this training should cover the capabilities, limitations, and key risks of general-purpose AI tools. The training should cover risk mitigation strategies, user responsibilities, applicable laws and regulations, and best practices for effective use (e.g., prompt engineering, evaluating AI outputs). Given the rapid evolution of AI capabilities and risks, VITA should annually reassess its AI training to ensure state employees remain up to date on emerging risks.

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.

Because not all employees access the internet regularly as part of their job, the state could consider limiting this requirement to state employees whose job responsibilities require them to use a computer or other electronic devices with access to the internet (e.g., tablets, phones) for a substantial portion of their work. For example, Texas’s AI training requirement applies only to employees who use a computer for at least 25 percent of their work duties.

VITA could also offer or facilitate voluntary, role-specific training on general-purpose AI tools. In JLARC’s survey, most agencies (89 percent) indicated that role-specific training would be helpful in supporting effective and responsible general-purpose AI use at their agency. For example, training tailored to agency IT staff could cover topics such as AI tool administrator security configurations to help mitigate risks and the responsible use of AI coding assistants. If pursued, VITA should consult with agencies to identify the roles and topics that would be most helpful for the role-specific training to cover.


4 Near-Term Opportunities and Barriers

The resolution for this study directed JLARC to evaluate both the opportunities and risks presented by AI tools. While Chapters 2 and 3 discuss how state agencies are using or planning to use AI tools and how to mitigate their risks, this chapter focuses on examples of near-term opportunities where AI tools, if used responsibly and effectively, could help agencies address longstanding operational challenges. This chapter also addresses several barriers identified by agencies and VITA that will likely constrain agencies' adoption in the near term, including challenges identifying a business case for AI use, data readiness issues, coordination and communication issues within VITA, and a lack of funding for AI projects.

AI capabilities have advanced rapidly in recent years, presenting near-term opportunities for state agencies. Benchmark data compiled by Stanford University's Institute for Human-Centered Artificial Intelligence (HAI) indicates that AI models now perform at or above human baseline levels on several key tasks, including language understanding, reading comprehension, and software engineering (Figure 4-1).

FIGURE 4-1

AI model performance has improved across a broad set of benchmark categories, with many approaching human baseline performance



SOURCE: Stanford Institute for Human-Centered Artificial Intelligence (HAI), 2026 AI Index Report.

NOTE: Stanford's HAI cautions that while benchmarks remain useful indicators of technical progress, they are vulnerable to reliability problems, flawed questions, saturation, and gaming.

AI use cases by federal agencies and other states indicate near-term opportunities for Virginia


Determining whether an AI tool should be used for a specific agency function requires evaluating multiple factors, including tool capabilities, the business case for adoption, and potential risks. Because these factors vary across functions, compiling an exhaustive list of potential use cases in state government is not feasible. However, reviewing established uses in the federal government and other states can help inform possible near-term opportunities for Virginia.

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

As AI capabilities have increased, both federal agencies and other states have started to adopt them to support their operations or services. The number of AI use cases deployed or piloted by U.S. federal agencies has increased considerably over the past decade, from six in 2016 to 340 in 2025, according to an inventory compiled by the U.S. Office of Management and Budget (Figure 4-2). Federal agencies reported an additional 31 use cases that were in the development or acquisition stages.

FIGURE 4-2

Federal agencies are increasingly using AI tools to support their operations or services (2025)



SOURCE: JLARC staff analysis of 2025 Federal Agency Artificial Intelligence Use Case Inventory, Individually Reported AI Use Cases, U.S. Office of Management and Budget. Accessed August 17, 2026.


NOTE: The U.S. Office of Management and Budget defines an AI use case as “the specific scenario in which AI is designed, developed, procured, or used to advance the execution of agencies’ missions and their delivery of programs and services, enhance decision making, or provide the public with a particular benefit.” Figure excludes common use cases involving commercial off-the-shelf products, three retired use cases, and use cases where federal agencies are not required to report to OMB, such as for national security or basic research purposes. Figure includes machine learning, natural language processing, and other classifications of artificial intelligence.

Federal AI use cases vary considerably but are concentrated in four areas: administrative functions, information technology, science, and law enforcement. These four areas comprised 231 of the 340 (68 percent) AI use cases that were initially piloted or deployed in 2025. Federal agencies were also piloting or deploying AI systems or tools for other functions, including service delivery, human resources, cybersecurity, and procurement and finance.

Federal agencies are also required to certify whether they use broadly available commercial AI products and services. In 2025, at least two-thirds of agencies reported using these tools, including for document summarization (87 percent), agency information retrieval (78 percent), and software code generation (69 percent) (Figure 4-3).

FIGURE 4-3

Federal agencies report using broadly available commercial AI products and services for various purposes (2025)



SOURCE: JLARC staff analysis of 2025 Federal Agency Artificial Intelligence Use Case Inventory, Consolidated Commercial Off-The-Shelf AI Use Case Inventory, U.S. Office of Management and Budget. Accessed August 17, 2026.

NOTE: The U.S. Office of Management and Budget defines an AI use case as “the specific scenario in which AI is designed, developed, procured, or used to advance the execution of agencies’ missions and their delivery of programs and services, enhance decision making, or provide the public with a particular benefit.” Figure shows use cases where at least two-thirds of reporting agencies indicated they were using AI for these tasks. It does not include all federal AI use cases. Figure excludes individual AI use cases reported by federal agencies (e.g., internally developed AI tools) and use cases where federal agencies are not required to report to OMB, such as for national security or basic research purposes.

No similar systematic data is available to understand the extent of AI use across states, but according to the National Association of State Chief Information Officers (NASCIO), there has been a “surge” in AI use cases in recent years. NASCIO’s 2025 State CIO survey indicated the five most common use cases among states were:

- chatbots and virtual assistants for citizen services and internal support;
- document generation, summarization, and policy analysis;

- code generation and developer productivity;
- content creation and workforce support; and
- analytics, translation, and specialized use cases (e.g., mapping and natural resources).

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

As discussed in Chapter 2, Virginia state agencies are already using or pursuing a wide range of use cases for AI technologies to streamline operations, reduce costs, improve staff productivity, improve data quality or accessibility, and assist with compliance, oversight, and fraud detection. Considering current technological capabilities and state needs, two key near-term opportunities for state use of AI include modernizing the state's dated IT systems and preventing and detecting errors in Supplemental Nutrition Assistance Program (SNAP) payments.

Migrating or replacing more than 850 state IT applications in need of modernization

AI coding assistants, if implemented with appropriate controls, could help state efforts to update outdated software and modernize agencies' legacy IT systems. As part of its oversight, VITA regularly assesses the condition of agencies' applications for cybersecurity vulnerabilities and operational fitness, identifying applications that need to be modernized or eliminated. As of August 2025, VITA had identified 865 Commonwealth IT applications that needed modernization, including 344 requiring urgent attention (sidebar).

As of June 2026, several state agencies were using or considering AI tools to help modernize their IT systems. In JLARC's survey of executive branch and independent agency IT leaders, six state agencies reported currently using AI assistance to modernize at least one system, and 18 agencies were exploring the approach (sidebar). No state agencies reported having fully completed a modernization project with AI assistance at the time of the survey.

In interviews and survey responses, several agencies cited using AI coding tools to help understand old IT systems without documentation, identify potential fixes to address IT security findings, and convert code into a more secure programming language. For example, one agency noted: "AI has proved to be invaluable in documenting 30-year-old code that has been modified hundreds of times over the years with fragmented documentation. It's also been a great resource in generating and testing new code." Another agency noted it is also exploring using AI coding assistance tools to help convert a 30-year-old legacy system away from a prohibited coding language to one that is more secure (sidebar).

Of the 1,733 state agency applications that VITA tracks, excluding out-of-scope agencies and higher education institutions, 800 (46 percent) are more than 10 years old, and 342 (20 percent) are more than 20 years old.

JLARC conducted a survey of senior IT leaders at executive branch and independent state agencies. JLARC received completed survey responses from all 80 agencies surveyed (100 percent response rate). See Appendix B for more information.

VITA's Enterprise Architecture Standard specifies 17 "prohibited" programming languages. Because these languages present substantial security risks, agencies are required to replace, retire, or migrate them to approved languages unless they receive an exemption.

At least one state has mandated that its state IT agency help agencies use generative AI tools to modernize their systems. In 2025, the Texas legislature passed a law directing its Department of Information Resources to help modernize or replace legacy IT systems with assistance from generative AI.

Given the volume of Virginia state IT applications needing modernization and their associated costs, VITA should prioritize supporting agencies' adoption of AI coding assistance tools (sidebar). This approach aligns with advancements in AI coding capabilities and their use in federal and state governments. As part of this effort, VITA should evaluate these tools and determine whether a statewide contract for one or more coding assistance tools would be appropriate. VITA should also develop specific guidance and requirements for agencies' use of coding assistance tools to support their responsible and effective use, including minimum training requirements for IT staff using these tools.

RECOMMENDATION 9

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

In January 2024, before AI coding assistants experienced substantial capability gains, **VITA-supported agencies identified \$477 million in funding needs for their most critical system and application modernization projects.** VITA has not conducted a similar assessment since, but it provides an indication of the magnitude of resources needed for IT modernization.

Reducing SNAP payment error rates

Another potential near-term strategic use of AI tools is to help reduce Virginia's SNAP payment errors, which could soon cost the state up to an estimated \$270 million annually under new federal requirements. Payment errors include both underpayments and overpayments that occur when either (1) an ineligible household receives benefits or (2) an eligible household receives the wrong amount of benefits. According to the Virginia Department of Social Services (VDSS), errors are often the result of either local social services department staff making data input or calculation errors or households providing inaccurate income or other required information.

Federal changes to SNAP under H.R. 1 (2025) require states to contribute to the cost of SNAP benefits if they have a payment error rate at or above 6 percent (sidebar). In January 2026, using the FY24 payment error rate (11.50 percent), VDSS estimated that Virginia would be required to pay \$270 million annually beginning in FY28, assuming the error rate remained the same.

Reducing the error rate could substantially lower the state's cost. For example, VDSS estimated that reducing the error rate to between 8.00 and 9.99 percent would lower Virginia's annual match requirement by \$90 million. The state would not need to pay any amount if its rate fell below 6 percent.

The required state match is tiered: 5 percent of benefit costs for error rates of 6.00 to 7.99 percent, 10 percent for error rates of 8.00 to 9.99 percent, and 15 percent for error rates of 10.00 percent or higher.

Several other states are using or planning to use AI technologies to help reduce SNAP payment errors, and the federal government has encouraged states' use of AI tools in SNAP administration. For example, Maryland deployed AI-supported training modules for caseworkers to practice with real-life scenarios as part of a broader strategy to improve SNAP payment accuracy. Additionally, the 2026 Florida legislature appropriated \$4 million to its Department of Children and Families to procure a vendor to help reduce its SNAP error rate. Florida required that the vendor “be capable of providing large-scale analyses of eligibility determinations using machine learning, identifying and correcting erroneous determinations, identifying root causes of erroneous determinations, and recommending operational improvements to avoid future errors.” In February 2025, the U.S. Department of Agriculture issued a memo expressing its support for states' pursuit of advanced automation technologies, including artificial intelligence, to improve SNAP administration as long as they are consistent with statutory and regulatory requirements.

As of September 2026, VDSS had deployed two AI tools to help prevent and detect SNAP errors, and the agency had identified other potential use cases. One tool, called the SNAP Policy Bot, is an AI-enabled information-retrieval chatbot intended to help local department of social services (LDSS) staff find relevant state and federal SNAP policy information without searching lengthy program manuals (sidebar). VDSS also recently deployed an AI-supported tool to identify potential errors in real time for LDSS before a benefit payment is certified. VDSS leadership reports the agency plans to pursue other AI use cases, such as a real-time interview assistant tool, to help reduce the state's payment errors.

Given the potential magnitude of state savings from reduced SNAP payment errors, VITA should prioritize its review and approval determination of any AI technologies proposed by VDSS to address error rates.

RECOMMENDATION 10

The Virginia Information Technologies Agency should prioritize its review and approval determination of artificial intelligence use cases identified by the Virginia Department of Social Services to reduce Supplemental Nutrition Assistance Program payment errors.

VITA should take several steps to increase agency awareness of potential business cases for AI use

As mentioned in Chapter 2, the leading reason why 38 percent of executive branch and independent state agencies had not pursued AI tools is that they “have not identified a business need or clear use case for AI.” While some agencies may lack a clear business case for AI adoption—often because of high costs, technical complexity, or risks—some agencies may not be pursuing AI tools because of a lack of awareness of how they could enhance operations or service delivery.

As of August 2026, VDSS's consolidated SNAP program manual was 714 pages.

VITA has taken some steps to increase agency awareness of potential use cases. For example, the agency hosts a monthly AI Community of Practice meeting to provide information and updates about AI use by supported agencies. A recent meeting, for example, highlighted a use case from the Department of Veterans Services.

VITA could further increase agency awareness of potential use cases by regularly publishing a list of pilot or deployed use cases that all agencies can review. The federal government and some other states publish a list of use cases, which could inform Virginia’s approach. These inventories, generally updated at least annually, provide information on a use case’s:

- vendor information, if applicable,
- sponsoring agency,
- service delivery area,
- purpose, and
- status, such as in pilot or deployed.

Similar to the federal government’s approach, VITA could have discretion to exclude certain use cases from public reporting for security or other compelling reasons.

RECOMMENDATION 11

The Virginia Information Technologies Agency should make available to state agencies a regularly updated list of piloted and deployed state agency artificial intelligence (AI) use cases to increase awareness of potential AI applications in state government that can support agency operations and service delivery.

“
A common theme with
VITA is there are things
available I didn’t know
about.
”

– Agency senior staff

State agencies cite several other barriers to adopting AI technologies

In interviews and survey responses, state agency staff cited several common barriers to their ability to adopt new AI tools. These include data readiness problems, funding constraints, and intra-agency coordination problems within VITA. Addressing these barriers could help further enable responsible and appropriate agency adoption of AI tools.

State data readiness challenges are a barrier to AI adoption and can raise risk of errors and bias

Data readiness is a key barrier to the implementation of AI tools cited by agencies and VITA staff. The majority of agencies currently using (77 percent) or pursuing (54 percent) function-specific AI cited data deficiencies (e.g., due to quality, labeling, or documentation) as a major or moderate constraint on their ability to pursue or implement AI.

Experts (e.g., Government Accountability Office, Organisation for Economic Co-operation and Development, Center for Democracy & Technology, NASCIO) emphasize the importance of adequate data preparation for producing reliable, unbiased AI outputs. AI models trained, fine-tuned, or recalibrated based on inaccurate, incomplete, or unrepresentative data may learn incorrect patterns, underperform for specific groups or data categories, or reflect historical inequities.

State agencies are subject to several requirements related to the readiness and quality of data used in AI systems and more broadly. VITA's Enterprise Architecture Standard for AI includes a variety of data readiness requirements, requiring AI system data to be high quality, validated, and standardized. Similarly, Virginia's Government Data Collection and Dissemination Practices Act requires agencies to maintain accurate, timely, and complete personal data to ensure fair benefit determinations. VITA staff noted that using biased or poor-quality data can also risk producing AI-assisted decisions that violate civil rights laws.

Causes of insufficient data readiness vary across agencies. Data deficiencies may result from poor data input processes, insufficient agency attention to data governance, inconsistent formatting practices within or across datasets, or unstructured data (e.g., write-in text fields) that is not in a format ready to analyze. The time and effort required to fix existing data issues also varies across agencies, according to VITA and agency staff.

Agencies can work with ODGA to improve their data readiness or purchase licenses for these tools from ODGA to perform scans and address issues themselves.

VITA's Office of Data Governance and Analytics (ODGA) offers data scanning tools that can help agencies improve the AI readiness of datasets, such as by flagging missing values, duplicates, and formatting inconsistencies, and helping build data glossaries. ODGA reported to JLARC staff that agencies' demand for scanning tools is increasing (sidebar).

Lack of coordination within VITA complicates state's AI adoption

VITA appears organizationally unprepared to be an effective and efficient partner in helping agencies adopt AI tools. VITA is organized into seven major functional areas, most of which have at least some role in the governance of AI systems and tools by VITA-supported agencies. These functional areas are further divided into sub-units, some of which agencies may need to interact with when they want to adopt or implement AI systems and tools.

As discussed in Chapter 2, agency AI adoption requires review and approval across VITA's various divisions. The number of divisions involved in the reviews can vary, based on factors such as cost, risk level, and whether the tool is internally developed or procured. Each division has its own requirements that agencies must meet before moving forward with their project or procurement.

Agencies expressed that poor coordination within VITA has directly affected their ability to adopt AI technologies, according to interviews, survey responses, and inter-agency meetings observed by JLARC. For example, of the 18 VITA-supported

agencies that reported seeking procurement of an AI tool over the past 12 months, 13 (72 percent) indicated that “coordinating required VITA reviews (e.g., legal, security, privacy)” was a major (five) or moderate (eight) challenge for their agency in their efforts to procure AI tools. The remaining either indicated it had been a minor challenge (three), or they did not know (two).

Examples of other issues cited by agencies include

- a lack of proactive updates from VITA about projects or issues,
- unclear processes that require pauses on projects in process,
- reliance on relationships outside of the process to expedite projects (e.g., going directly to the chief information officer [CIO]), and
- unclear navigation within VITA for questions and issues by either agencies or VITA staff assigned to support agencies (e.g., a single point of contact for AI questions).

Currently, VITA leadership lacks insight into when and why coordination delays, bottlenecks, and points of confusion occur across the agency’s various review and approval processes for AI technologies and use cases. Without this information, leadership cannot effectively streamline the process or address agency concerns. VITA should identify where delays and confusion are the most problematic. (Recommendation 5 in Chapter 2 related to documenting the complete review process for AI technologies and use cases could serve as a starting point.)

To assist with this effort, VITA’s new CIO could consider establishing one or more positions responsible for coordinating VITA’s support and oversight of agency AI projects, as well as ensuring that VITA’s AI guidance and requirements are clear. No single position currently holds this overall responsibility, which can make internal accountability difficult because of the multiple units and staff involved. Any such position would need sufficient authority to identify and address communication problems and bottlenecks.

RECOMMENDATION 12

The Virginia Information Technologies Agency (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.

The lack of coordination and confusion surrounding VITA’s approval of AI projects is part of a systemic issue that has existed for years within VITA. VITA staff have repeatedly expressed concerns regarding internal coordination in internal culture surveys. For example, VITA’s 2025 internal culture survey highlighted unfavorable staff perspectives regarding the agency’s cross-functional collaboration, timely decision making, and work processes. Only 29 percent reported that VITA’s work processes were highly efficient. Similar sentiments regarding VITA’s siloed environment were

“ There are too many silos, especially in project management. We need to do a better job of laying out project needs. I would benefit from having someone who keeps us on track for each project. ”

– VITA staff

expressed in the 2022 culture survey. While VITA staff reported good inter-team collaboration, they were less positive about collaborations across functional areas.

More than two-thirds of agencies pursuing or using AI tools cite funding constraints as a barrier to adopting them

At least some additional funding would likely be needed to increase AI adoption substantially across agencies, according to VITA staff and agencies. Seventy-two percent of agencies that had been using, pursuing, or considering AI systems as of June 2026 cited “insufficient funding for new AI systems or tools” as being a major (40 percent) or moderate (32 percent) constraint on their ability to pursue or use them. Funding constraints appear to be affecting both large and small agencies’ ability to adopt AI tools.

No comprehensive data either documents agencies’ current spending on AI tools or forecasts agencies’ expenditures on AI tools in the next several years. The primary general-purpose AI tool available to VITA-supported agencies (Copilot Chat) is included in the cost of the broader Microsoft license and is not broken out separately. Agencies use existing resources to purchase or develop AI tools, as there is no separate budget item for their development, acquisition, or use. Agencies’ costs for AI deployment and use will vary based on many factors, including model costs and utilization rates. For example, the cost of a chatbot designed to answer citizen questions will largely depend on usage volume, and this use can be challenging to predict initially. In interviews, VITA leadership reported they are considering various approaches to charging state agencies for AI tools.

5 Mitigating AI-Related Risks to the Public

Artificial intelligence (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 technological and policy complexities.

JLARC staff conducted extensive interviews with experts and reviewed various reports to identify the challenges in regulating AI harms and to help inform these legislative efforts. Interviews were conducted with stakeholders, including industry representatives from Google, Microsoft, Amazon, and Anthropic; Virginia legislators; staff of the Division of Legislative Services (DLS), the Joint Commission on Technology and Science (JCOTS), the Office of the Attorney General (OAG), and the State Corporation Commission; legislative staff in other states; and national experts. JLARC also reviewed other states' legislation and key reports on the benefits and risks of AI tools and risk mitigation strategies.

AI adoption is rising, increasing public's exposure to both benefits and risks

Public use of AI is growing, allowing both businesses and consumers to take advantage of its potential benefits. However, as AI use increases, exposure to potential harms also becomes more widespread. State legislatures are acting to protect the public from potential AI harms, but their approaches vary widely and some legislative attempts have encountered obstacles.

AI technologies are increasingly being used by the public and businesses


A growing share of Americans are using AI tools. According to a 2026 Pew Research Center survey, about half of U.S. adults (49 percent) have used AI chatbots, such as ChatGPT, Gemini, or Copilot, in 2026, up from a third in 2024 (Figure 5-1). AI chatbot use among teenagers is even higher, with 64 percent of teenagers reportedly using AI, and 28 percent reporting they use chatbots at least daily, according to a 2025 Pew Research Center survey (Figure 5-2).

Adults reported using chatbots for a variety of purposes, including searching for information, editing photos or videos, receiving medical advice, and seeking emotional support. For example, nearly 30 percent of adults reported using AI tools monthly to

look up health information and advice, according to a Kaiser Family Foundation survey. Teenagers' top three reasons for using AI chatbots were information searches, schoolwork, and entertainment, according to a 2026 Pew Research Center article.

FIGURE 5-1

A growing share of American adults are using AI chatbots




SOURCE: Pew Research Center surveys from 2024 and 2026.

NOTE: 2024 survey: N=5,410. 2026 survey: N=5,119. As AI becomes increasingly integrated into existing software and search engines, the share of U.S. adults who have used an AI chatbot may be harder to measure, since AI is now embedded in many tools not typically thought of as AI.

FIGURE 5-2

Over a quarter of American teens report using AI chatbots at least daily



SOURCE: 2025 Pew Research Center survey.

NOTE: The "Do not use AI chatbots" category includes both teens who reported not using chatbots and those who had not heard of AI chatbots. N=1,458 pairs (one teenager and one parent per pair).

Businesses are also using AI, although adoption rates appear to vary by business size and sector. U.S. Census Bureau data shows that businesses reported using AI for marketing, business development, research and development, and information technology. The same data indicates that businesses with over 250 employees are using AI tools at a higher rate than smaller businesses. Additionally, the information, finance, and insurance sectors reported using AI tools more than other sectors, such as manufacturing and retail trade. Around 23 percent of Virginia businesses reported using AI, according to the U.S. Census Bureau's Business Trends and Outlook Survey data from late July 2026.

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

Experts indicate that AI tools can cause or contribute to public harm when used irresponsibly or misused. The MIT AI Risk Repository has documented over 1,700 risks as of July 2026, including exposure to toxic content, discrimination against protected groups, and the development of cyber weapons. Additionally, the 2026 International AI Safety Report, produced by an expert panel representing over 30 countries and international organizations, identified various risks and grouped them into categories:

- malicious use from bad actors (e.g., producing AI-generated content for crimes, executing cyberattacks);
- malfunctions from AI systems (e.g., producing false information, giving misleading medical advice); and
- systemic risks from widespread deployment (e.g., impacting labor markets, influencing critical thinking and decision making).

Major AI developers also acknowledge that AI carries risks of potential harm to the public. Anthropic (the developer of Claude) identifies “high-risk use cases” in its usage policy, such as using AI for medical, insurance, financial, employment, and housing decisions, and its policy prohibits using it for various harmful activities, such as humiliating individuals, creating deceptive or misleading information about laws, and creating falsified government documents. OpenAI (the developer of ChatGPT) has a policy that also prohibits using its models to commit violent acts against oneself or others, or to offer advice in regulated professions without a license.


Limited systematic data exists on realized harms in the U.S. or Virginia, but evidence indicates that AI is contributing to harmful activity. In its 2025 Internet Crime Report, the FBI reported that AI is making it easier for criminals to carry out fraud and other criminal activity. The FBI reported receiving more than 22,000 AI-related complaints nationally, with associated losses exceeding \$893 million. In Virginia, total reported financial losses to victims of AI-related crimes were approximately \$22.8 million. Public reports also highlight other realized harms, including AI chatbots engaging in exchanges that preceded self-harm or violence and recent cybersecurity incidents.

State legislatures are increasingly considering and enacting legislation to address AI with mixed success

States are increasingly enacting AI-related legislation to address different aspects of AI, including to mitigate its harms to the public. According to the Stanford Institute for Human-Centered AI, state legislatures enacted more than three times as many AI-related bills in 2025 compared with 2023, and more than five times as many as in 2022 (Figure 5-3).

FIGURE 5-3

States have increasingly enacted AI legislation, according to the Stanford Institute for Human-Centered AI



SOURCE: The Stanford Institute for Human-Centered AI (HAI) 2026 AI Index Report.

NOTE: According to Stanford's HAI report, bills were identified by searching "artificial intelligence" into state-level databases and then supplemented with bills identified on the Multistate AI Legislation Tracker.

Automated decision-making technology, under Colorado law (SB 26-189), refers to any system that processes personal data to generate outputs—such as scores, rankings, or recommendations—that inform or automate decisions about an individual. This scope includes advanced AI models as well as simpler, rules-based computer systems.

Synthetic content refers to digital media, such as text, audio, images, or video, that have been significantly altered or generated by AI algorithms (NIST). A common type of synthetic content is a deepfake, which uses AI to create realistic media depicting statements or actions that never actually occurred (Cybersecurity & Infrastructure Security Agency).

States are taking different approaches to mitigate potential harms caused by using AI tools. Some states, such as Colorado, have passed bills that create transparency and disclosure requirements for automated decision-making technology or reporting requirements for large "frontier model" developers (New York) (sidebar). Alternatively, some states have taken a more targeted approach, focusing on regulating AI use within a specific sector or context, like requiring disclosures for synthetic content during election campaigns (North Dakota) or barring chatbots from using official nursing titles, such as registered nurse (Oregon) (sidebar). Approaches also vary on whether bills seek to regulate AI technology *itself* versus the *use* of the technology. Washington's HB 2225, which specifically regulates AI companion chatbots, is an example of regulating the technology itself. Illinois's HB 1806, which limits AI use in mental healthcare, is an example of regulating the use of the technology.

As of August 2026, the most common AI-related legislation introduced by states addresses AI-generated deepfakes, consumer protections, public-sector AI use, child safety, and specific AI technologies or use cases, such as automated decision-making technology, generative AI, frontier models, and chatbots. AI-related legislation has also addressed AI use in certain sectors, such as healthcare, employment, and elections.

AI-related legislation has encountered various difficulties, however, including confusion over its scope, legal challenges, and recent federal actions. As discussed later in this chapter, a technical definition in Colorado’s original algorithmic discrimination law contributed to confusion between the legislative and executive branches over which technology systems it covered (sidebar). Courts have blocked two California election deepfake laws (AB 2839, AB 2655) on the grounds of free speech and federal preemption of state law (through Section 230 of the Communications Decency Act, sidebar). Lastly, a December 2025 federal executive order set federal priorities for AI policy and penalties for states that pass legislation contrary to these priorities. These potential consequences have reportedly made state lawmakers, including those in Virginia, more hesitant to pass new AI legislation.

Algorithmic discrimination, as defined in Colorado’s original enacted legislation (SB 24-205), is “any condition in which the use of an artificial intelligence system results in an unlawful differential treatment or impact that disfavors an individual or group of individuals on the basis of their actual or perceived age, color, disability, ethnicity, genetic information, limited proficiency in the English language, national origin, race, religion, reproductive health, sex, veteran status, or other classification protected under the laws of this state or federal law.”

AI can be challenging to legislate for several reasons

Regulating AI effectively can be difficult because of its rapid evolution, technological complexity, potential benefits and risks, use across various sectors, and overlap with existing law. Identifying potential challenges upfront helps ensure new laws are necessary and effective; however, addressing them requires a deep understanding of technology law and policy, and in some cases, other types of specialized legal or policy expertise (Table 5-1).

Existing laws and regulations may already address some AI-related 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, regulating the *underlying conduct, practices, or decisions*, rather than the *means* by which a harm may be occurring (e.g., a human, an automated system, or another technology). Examples of technology-neutral statutory provisions relevant to *businesses* using AI in ways that may cause harm include those pertaining to:

- consumer deception (§ 59.1-200);
- discriminatory employment practices (§ 2.2-3905); and
- discriminatory credit decisions (§ 6.2-501).

Examples relevant to *individuals* using AI tools in ways that may cause harm include those addressing:

- identity theft (§ 18.2-186.3);

Section 230 of the federal Communications Decency Act states that providers and users of “interactive computer services” cannot be held liable as the publisher for content provided by another person. This provision of federal law will be discussed later in this chapter, as it relates to liability.

- stalking (§ 18.2-60.3); and
- the dissemination of non-consensual intimate videos or imagery (§ 18.2-386.2).

TABLE 5-1

AI legislation presents potential challenges and questions that need to be addressed

Potential challenge	Example questions needing to be addressed
Existing laws or regulations may already address harms <i>Existing laws and regulations may already provide protection against some harms.</i>	• Do any existing laws and/or regulations currently provide protection against the harm? • If so, has this protection been tested or litigated, or is this a theoretical protection? • Does this bill conflict with existing law or regulations? • If not, what gaps exist in current laws and regulations where protection may be warranted?
Constitutional conflicts <i>State laws that conflict with constitutional provisions risk being struck down.</i>	• Could this bill be challenged as unconstitutional? • Have similar laws in other states faced legal challenges? If so, on what grounds?
Unclear technical definitions <i>Key terms can be over-inclusive or become quickly outdated.</i>	• What, if any, technology does the bill seek to regulate? Or should the language be technology neutral? • Are technical terms defined clearly enough to understand the specific technology or conduct governed by the bill? • Will the definition hold as the technology changes, and does it align with other states' laws?
Assigning liability <i>Multiple actors involved; modified or autonomous systems can further blur responsibility.</i>	• To what extent can the harm originate from decisions made by the developer, deployer, end user, or other party? • Which parties can prevent the harm and should be held responsible? • What happens if an open-weight model is involved and/or an agentic AI tool is used?
Difficult compliance requirements <i>Requirements may be difficult to meet with current technology or costly/burdensome to carry out.</i>	• Can the evidentiary and compliance requirements be met with current technology? At what cost? • Does the requirement fall disproportionately on small firms? • Which entity will be responsible for verifying or enforcing compliance, and do they have the capacity for this responsibility?
Harms depend on context <i>The same AI system can be beneficial or harmful depending on how it is used. AI uses carry different risk profiles and differ sector-by-sector.</i>	• Does the bill seek to regulate the technology itself or the use of the technology in a specific sector or context? • What sector-specific expertise may be needed to ensure the legislation works and is needed in practice? • Could the bill restrict potentially beneficial uses of AI systems in the policy area? • Are the burdens of this bill proportionate to the risk it is seeking to regulate?
Limited data on realized harms <i>Evidence is often hypothetical, self-reported, or inaccessible.</i>	• What evidence exists to indicate the extent to which the harm is occurring or is likely to occur? Is this evidence reliable?

SOURCE: JLARC staff analysis.

Existing regulations may also provide protection against potential AI-related harms. For example, the State Corporation Commission’s (SCC) Bureau of Insurance regulates how insurers make and communicate insurance decisions using existing statutory authority and rules in the Administrative Code, including the requirement that insurers must provide a “reasonable written explanation of the basis for any claim denial.” In addition, laws and regulations that establish requirements for the practice of certain licensed professions include provisions that may address potential harms caused by AI in those contexts. Thus, in some areas, new laws and regulations may not be needed to address potential harms from this new technology.

While existing laws and regulations may address some AI-related harms, legal gaps and evidentiary challenges may remain. For instance, although a given law prohibiting unfair discrimination in a particular context may be technology neutral, detecting and proving discrimination from an AI system may be difficult in practice, given the complex nature of these systems (discussed later in this chapter).

Legislators have several means to determine whether existing law adequately protects Virginia businesses and citizens from AI-related harms or whether new legislation may be needed. The OAG is required to provide legal advice to legislators upon request, and, according to staff, this advice can include an assessment of whether proposed legislation would be duplicative of existing law or would potentially provide additional protection from harm. The OAG also reviews each proposed bill to determine, among other things, whether it conflicts with existing laws and will advise patrons of potential conflicts. OAG staff noted that legislators may request this advice proactively, as needed. Similarly, DLS is required to “prepare summaries of existing laws affected by proposed legislation,” upon request of any member of the General Assembly.

Constitutional conflicts can invalidate AI-related legislation

The United States Constitution is the highest legal authority in the U.S., taking precedence over federal and state laws when there is a conflict. Similarly, state-level constitutions also take precedence over state-level laws. Therefore, proposed legislation and regulations on AI technology could potentially conflict with federal or state constitutional protections, such as free speech, privacy, and due process, leading them to be ruled unconstitutional.

Some AI legislation has already been found to conflict with constitutional protections. For instance, courts have blocked or struck down at least two election deepfake laws on First Amendment grounds, including California’s AB 2839 (2024) and Hawaii Act 191 (2024).

Proposed legislation seeking to regulate AI harms may encounter legal challenges under other constitutional provisions, including the Supremacy Clause (Article VI, Clause 2), which articulates that federal laws take precedence over state laws, and the Commerce Clause (Article I, Section 8, Clause 3), which gives Congress the power to regulate interstate commerce.

Unclear technical definitions, liability, and compliance mechanisms can cause AI-related legislation to be ineffective or unclear

Crafting effective AI legislation is complex because of the potential need to define specific technologies, assign liability across AI development, deployment, and use contexts, and determine practical and appropriate compliance mechanisms. If these areas are not thoroughly addressed, legislation risks being confusing, ineffective, or overly burdensome.

Unclear technical definitions can undermine AI-related legislation

There is no universally accepted definition of AI, which can result in legislative definitions that are unclear or overly broad. For example, an early proposed definition for AI systems in the European Union’s AI Act listed techniques such as machine learning, logic-based approaches, and statistical methods, which was so broad it could capture nearly all computer programs.

As mentioned in Chapter 1, the Organisation for Economic Co-operation and Development (OECD) released an updated, general definition in 2024, defining the term “AI system” as

a machine-based system that, for explicit or implicit objectives, infers, from the input it receives, how to generate outputs such as predictions, content, recommendations, or decisions that can influence physical or virtual environments. Different AI systems vary in their levels of autonomy and adaptiveness after deployment.

This definition is widely used, though OECD cautions its broad definition should be tailored to specific contexts when used in practice (sidebar). Failing to narrow the definition can create legal challenges. For example, Colorado’s SB 24-205 (2024), enacted to protect consumers from algorithmic discrimination, used the OECD’s definition of “AI systems.” This broad definition contributed to disagreement between Colorado’s executive and legislative branches 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.

OECD notes in their March 2024 memo on their updated definition of AI systems that their definition is inclusive, and since some techniques underlying AI systems have different risk profiles, “when applied in practice, additional criteria may be needed to narrow or otherwise tailor the definition when used in a specific context.”

Liability for AI-related harms remains an open question

Assigning liability is another complicated aspect of AI-related legislation. AI systems often involve different parties, and these parties may include:

- developers, who build and train the model;
- deployers, who buy, adapt, and deploy the model for a specific purpose; and
- end users, who interact with the system.

Determining responsibility for AI harms can be difficult because multiple parties can contribute to the harm. Harms may also be caused, in part, by design choices, such as fine tuning, 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.

Agentic and open-weight AI systems can also complicate liability questions. Agentic AI systems, which act with varying degrees of independence, could become another party in the chain of potential liability. An open weight model can also complicate matters, as third parties can download and modify the model’s weights, potentially altering the system from the original developer’s design (sidebar).

States are still evaluating how to best assign liability, and different approaches remain untested. One approach considered at both the state and federal levels relies on product liability law, which places responsibility on manufacturers and sellers for defective products. For example, Illinois introduced legislation (SB 3502) in 2026 that would have allowed product liability claims against AI developers for defective designs, inadequate warnings, or breach of warranty. Under this proposal, deployers would be liable if they substantially altered the system or intentionally misused it in a way that causes harm.

Liability for AI-related harms is also complicated by federal law. Section 230 of the Communications Decency Act states that providers and users of “interactive computer services” cannot be held liable as the publisher of user-provided content. While this provision has historically shielded online platforms from liability for user-generated content, the judiciary has not yet determined how it applies to AI-generated outputs. Until its relevance to AI systems is clarified by Congress or federal courts, state law attempting to assign liability to AI companies could be ineffective because of Section 230 immunity.

Some AI-related bills may include compliance requirements that are difficult to meet given current technological limitations

Current technological limitations may make some compliance requirements in AI-related legislation difficult to fully meet. For example, the original version of Colorado’s SB 24-205 would have required deployers to provide a statement disclosing the principal reason or reasons for an adverse “consequential decision,” such as one related to employment or housing, when an AI system made or substantially influenced that decision. Fully complying with such a mandate may be challenging because the internal decision-making of many advanced AI systems is often not fully explainable, even to the developers of the system (sidebar). While organizations may be able to produce an explanation to satisfy the legal requirement, experts caution that these explanations may not accurately reflect how the AI system reached its decision.

While full explainability remains technologically challenging, states have enacted measures to achieve transparency and safety goals. For instance, some states have required frontier AI developers to publish AI safety frameworks, report critical safety incidents (discussed further below), and conduct independent third-party audits of their compliance with state-level AI regulations.

With open-weight models, the developer releases the model’s weights for public download and modification but does not release its training data. Anyone who downloads an open-weight model can adjust the model’s parameters for their own purposes.

Fully explaining how modern AI systems produce their outcomes can be difficult or impossible, as it is unclear exactly how many AI systems reach a certain output. Because of the technology’s complexity, this decision-making process is often opaque not only to end users but to the developers themselves.

Current technological limitations may also make legislation seeking to address AI-related harm difficult or costly to enforce, such as when enforcement depends on proving that content is AI-generated. In 2025, the Virginia General Assembly passed legislation that would have created a separate Class 1 misdemeanor for using synthetic digital content to commit fraud, contingent on re-enactment. In a legislatively mandated January 2026 report, a workgroup convened by the Virginia OAG concluded that “establishing that fraud was committed through the use of synthetic digital media presents significant evidentiary hurdles.” The workgroup reported that, according to experts, “reliably determining whether digital media is synthetic remains difficult and, in many cases, uncertain.” It added that even when such a determination might be possible, it “would almost certainly require expert testimony,” adding expenses and delays to prosecutions. The legislation was not re-enacted during the 2026 General Assembly session.

Risk often depends on context in which AI system is used, making regulation challenging

AI tools are often “dual use,” meaning the underlying technology could be used in both beneficial and harmful ways. For example, an AI tool used to generate new code could also be used to write malware. Because beneficial and harmful uses are possible from the same underlying capability, broad efforts to mitigate potential harms could also limit potential beneficial uses.

Additionally, the same AI tool can have different risk levels depending on how it is used or the context in which it is used. A general-purpose chatbot presents minimal risk when being used to help draft or edit a meeting agenda but can create significant risk when providing medical advice. The magnitude of possible uses and contexts makes it difficult to anticipate and identify every use case and the associated risks, or test how an AI tool might behave in different contexts.

Systematic data on realized AI-related harms is limited, making it difficult to gauge their prevalence and need for regulation

Systematic evidence of harms enabled or caused by AI systems is currently limited, making it difficult to gauge their breadth, prevalence, and severity. Without this insight, both over- and under-regulation of AI is possible. Reliable data on realized harms can help ensure legislative efforts are targeted and proportionate.

Some states have passed legislation that should provide more information on AI-related harms. Since September 2025, three states (California [SB 53], New York [A6453A], and Illinois [SB 315]) have enacted laws regulating frontier AI models, which are the largest, most capable general-purpose AI models. Each requires frontier AI model developers to report critical safety incidents to a relevant government entity such as an Emergency Services agency and Attorney General. Information sharing may give governments early insight into emerging risks and better data on harms as

they occur. As of August 2026, only California’s law was in effect, so the impact of these requirements is unclear.

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

Developing effective AI legislation can require a deep understanding of various legal, technological, and policy challenges, including those noted previously. Two improvements would strengthen the information and expertise available to assist legislators during the bill drafting and deliberation process. These include ensuring DLS has access to AI-specific technology law expertise and providing legislators with routine, updated information on AI.

As discussed previously, AI-related legislation may also raise complex legal questions beyond the technology, such as those involving tort or constitutional law. To advise legislators effectively, DLS needs to continue to have attorneys on staff, or access to outside attorneys, who can provide appropriate counsel on these issues as well.

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

DLS is the nonpartisan legal and general research agency of the General Assembly. The division is statutorily tasked with drafting bills and advising legislators “as to the constitutionality or probable legal effect of proposed legislation” upon request. DLS serves a key role in helping legislators by drafting legislation that achieves its goals, avoids unintended consequences, and does not encounter foreseeable complications.

To help legislators navigate the challenges posed by AI-related legislation, DLS needs staff with specialized technology law expertise, particularly related to AI, or access to such expertise. For example, advising on the assignment of liability requires an in-depth understanding of the actors involved across AI development, deployment, and use contexts, such as the model developer, deployer, and end user. Similarly, advising on appropriate compliance-related requirements, such as explainability, requires an understanding of whether it is possible to meet such requirements with existing technological capabilities and limitations.

Specialized technology law expertise could also be a resource to help legislators determine how AI-related legislation should be handled during the session, as needed. Proposed legislation may often raise both policy and technology issues, requiring a determination of whether it would be better assigned to a policy committee, such as the House Health and Human Services Committee, or a technology committee such as the House Communications, Technology, and Innovation Committee. For example, a bill seeking to mitigate harms associated with the use of AI chatbots for mental health

services without causing unintended consequences for users may raise both technology and behavioral health issues.

Having staff with specialized technology law expertise or access to such expertise could also help legislators determine when proposed AI-related legislation should be referred to a legislative commission for a more in-depth review and which commission would be most appropriate. For example, JCOTS may be best suited to address legislation seeking to regulate the *development or deployment* of AI technologies, but other commissions, such as the Joint Commission on Health Care, may be better able to evaluate legislation seeking to regulate the *use* of AI technologies in a particular policy context.

In interviews, both legislators and DLS staff indicated the need for technology law expertise at DLS. As of September 2026, DLS does not have on staff, or access to, an attorney specializing in technology law to advise the General Assembly on AI-related legislation.

DLS should determine how to best obtain specialized technology law expertise, going beyond a basic understanding of technology terms and concepts, to effectively advise legislators on AI-related legislation on an ongoing basis. This could involve hiring one or more attorneys specializing in technology law, and who have specific expertise in AI and public policy. Attorneys with this expertise may command high salaries; therefore, DLS could consider whether partnering with a university with this expertise (some law schools offer advanced credentials in technology law and policy) or retaining outside counsel is more feasible. If DLS ultimately decides to obtain outside expertise in technology law, it should still designate a lead attorney on staff as the primary AI point of contact. This attorney should have enough understanding of the technological, legal, and policy issues around AI to determine when special expertise is necessary and to provide guidance, when needed, on directing AI-related legislation to the appropriate legislative committee or commission.

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 artificial intelligence legislation.

DLS's new attorney(s) specializing in technology law, or externally acquired expertise, should work with other DLS attorneys to provide support when AI-related legislation arises across policy areas. This specialized expertise, whether hired in-house or acquired externally, could also help other DLS attorneys build a baseline knowledge of AI concepts, such as through informal trainings, so that all DLS attorneys are better equipped to identify and address potential problems in AI legislation.

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

JCOTS is a legislative commission statutorily charged with advising Virginia leaders, agencies, and institutions on science and technology matters. As such, the commission is well positioned to keep legislators informed of key AI concepts and terminologies, important developments in AI technologies, and benefits and harms from AI use. This could be particularly helpful given the lack of systematic and defensible data available to legislators (as discussed previously in this chapter).

JCOTS should be directed to provide the General Assembly with a legislator's guide on AI, similar to the one JLARC produces for the Virginia Retirement System. This short, annually updated guide should cover topics like important and relevant AI terminology, key legislative considerations, and potential, actualized, and emerging benefits and harms from AI use. Other useful information could include a summary of recently enacted AI legislation in Virginia and elsewhere and recommendations of reputable resources for legislators if they have additional questions about AI. The goal should be a clear, concise reference for legislators to use as they develop or deliberate AI legislation. JCOTS could partner with the Virginia Academy of Science, Engineering, and Medicine to assist with researching and drafting the guide if outside expertise is needed.

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.

Appendix A: Study resolution

Artificial Intelligence in the Commonwealth

Authorized by the Commission on December 16, 2025

WHEREAS, artificial intelligence (AI) can enhance the delivery of government services, reshape private sector business processes and productivity, and significantly impact citizens' day-to-day lives; and

WHEREAS, while AI brings opportunities for improving productivity and government services, it also brings risks such as supplanting human actions and judgment, perpetuating discrimination and inequality, creating security risks such as deepfakes and other illegal activity, and undermining public trust in government institutions; and

WHEREAS, Virginia has started taking actions to oversee and regulate the appropriate use of AI in both the public and private sector, including Executive Order 30 in 2024, the Virginia Information Technologies Agency's (VITA) issuance of AI IT security standards and an AI registry for state agencies, and the Joint Commission on Technology and Science's (JCOTS) reviews of some specific aspects of AI; and

WHEREAS, legislation is being considered and/or adopted in other states and Virginia to appropriately regulate AI to protect citizens while also harnessing opportunities to enhance productivity, processes, and prosperity; now, therefore, be it

RESOLVED by the Joint Legislative Audit and Review Commission that staff be directed to review the use, governance, regulation, and oversight of AI in the Commonwealth. In conducting its review, staff shall: (i) determine how and to what extent AI is currently being used and considered in Virginia state government [state agencies and higher education] and the opportunities and risks these uses present; (ii) determine whether state agency staff have the appropriate skills to effectively operationalize the use of AI; (iii) evaluate whether state policies and processes regarding the use of AI in state government enable innovation while also mitigating the risks to state personnel and the public; (iv) identify options for a governance structure and process the state could adopt to ensure continuing rigorous oversight of AI as it evolves; (v) review legislation pertaining to AI that has been considered and passed in other states; (vi) identify areas outside of state government that are likely to be substantially affected by AI that could require state government regulation; and (vii) review how data center capacity and other factors could impact the cost and availability of AI.

JLARC shall make recommendations as necessary and review other issues as warranted.

All agencies of the Commonwealth, including the Virginia Information Technologies Agency, the Office of Regulatory Management, and higher education institutions shall provide assistance, information, and data to JLARC for this study, upon request. JLARC staff shall have access to all information in the possession of agencies pursuant to § 30-59 and § 30-69 of the Code of Virginia. No

provision of the Code of Virginia shall be interpreted as limiting or restricting the access of JLARC staff to information pursuant to its statutory authority. JLARC will coordinate as necessary and appropriate with JCOTS.

Appendix B: Research activities and methods

Key research activities JLARC performed for this study included:

- structured interviews with leadership and staff of the Virginia Information Technologies Agency (VITA) and other state agencies, industry stakeholders, and subject-matter experts in Virginia and nationally;
- a statewide survey of executive branch and independent agency IT leadership;
- analyses of state employee general-purpose AI usage data;
- analyses of federal AI use case data; and
- reviews of relevant documentation, including national standards and reports on AI, other states' AI-related policies and legislation, and laws, regulations, and policies relevant to AI use and governance in Virginia.

Structured interviews

Structured interviews were a key research method for this report. JLARC conducted around 70 interviews. Key interviewees included:

- VITA leadership and staff;
- IT leadership and staff of state agencies;
- industry representatives; and
- stakeholders and subject-matter experts in Virginia and nationally.

VITA leadership and staff

JLARC conducted 16 interviews with VITA leadership and staff across various VITA divisions. Topics varied across interviews but were primarily designed to understand VITA's oversight and administrative role in state agency AI adoption, understand the state's current policies and processes for managing AI-related risks, and identify potential opportunities to improve the guidance and tools VITA provides to agencies. Interview topics included:

- VITA's overall role in AI oversight, current priorities, and the most significant challenges to supporting responsible AI adoption across agencies;
- the AI registry's purpose and approval workflow, planned changes to the registry platform, and how "high-risk" AI systems are determined for annual audit requirements;
- barriers to agency AI readiness, risks associated with agencies developing their own AI systems, and the development of VITA's AI governance framework, utilization policy, and AI literacy training for state employees;
- the processes and risk reviews applied to AI procurements, including the state's cloud services security review process, and ongoing efforts to negotiate contract terms for AI-related cloud services with vendors;
- IT security-related reviews and mechanisms for mitigating AI-related risks associated with state agency and state employee use of AI tools;

- agency data readiness for AI, common data quality challenges, and agency compliance with data governance reporting requirements; and
- cybersecurity risks posed by AI, VITA's own use of AI to support security operations, and risks specific to agentic AI tools.

Staff of other state agencies

JLARC also interviewed senior IT staff, including chief information officers or comparable positions, at 13 executive branch and independent state agencies including:

- Department of Environmental Quality;
- Department of Forestry;
- Department of Medical Assistance Services;
- Department of Motor Vehicles;
- Department of Social Services;
- Department of Veterans Services;
- State Corporation Commission;
- Virginia Department of Health;
- Virginia Department of Transportation;
- Virginia Economic Development Partnership;
- Virginia Employment Commission;
- Virginia Workers' Compensation Commission; and
- Virginia Works.

Agencies were selected to reflect a range of sizes, mission areas, and levels of AI activity. Interviews varied but generally covered:

- the agency's overall approach to AI, how it identifies AI opportunities, and the status of any registered use cases with VITA;
- the greatest near-term opportunities the agency sees for AI to improve its operations or service delivery;
- barriers the agency has encountered in adopting AI, including staffing, data quality, state requirements, and procurement;
- whether the state's AI policies and processes strike an appropriate balance between enabling innovation and mitigating risk, including the clarity and usefulness of the AI registry process;
- AI-related risks of greatest concern to the agency, including risks specific to its mission area;
- staff use of general-purpose generative AI tools; and
- the agency's internal guidance and training for staff use of generative AI, and what additional state guidance or support would be most helpful.

As part of JLARC's research for Chapter 5, JLARC staff interviewed staff of the Division of Legislative Services, the Joint Commission on Technology and Science, and the Office of the Attorney General.

Industry stakeholders

JLARC staff conducted interviews with industry stakeholders including representatives from Amazon, Anthropic, Google, Microsoft, and the Northern Virginia Technology Council (NVTC). Interview topics generally focused on opportunities for how legislators could mitigate AI-related harms without causing undue burdens or unintended consequences and key factors affecting the cost and availability of AI.

Other stakeholders and subject-matter experts in Virginia and nationally

JLARC staff also interviewed Virginia and national subject-matter and company experts, including from:

- Atos;
- Center for AI and Digital Policy;
- Center for Democracy & Technology;
- Georgetown University's Center for Security and Emerging Technology;
- Institute for Trustworthy AI in Law & Society;
- National Association of State Chief Information Officers;
- National Conference of State Legislatures;
- New York University's Center on Technology Policy;
- RAND Corporation;
- SAIC;
- Stanford University's Institute for Human-Centered Artificial Intelligence;
- Stanford University's RegLab;
- University of Virginia's School of Data Science;
- Virginia Academy of Science, Engineering, and Medicine; and
- William & Mary Law School's Digital Democracy Lab.

These interviews were used to gather interviewees' perspectives on a variety of topics, including opportunities and barriers to AI adoption in state governments, effective governance approaches, and actions taken in other states or jurisdictions for oversight or guidance on adoption. Other topics covered included benefits and risks associated with AI tools, different state approaches to regulating AI, how existing laws may protect against harms from AI systems, and the key factors impacting the cost and availability of AI.

Survey

JLARC staff administered an electronic survey from June 2026 to August 2026 to senior IT leadership at executive branch and independent state agencies. This survey was for senior leaders primarily responsible for technology-related decisions at executive branch and independent state agencies, such

as chief information officers or equivalent senior IT leaders. For agencies without dedicated IT staff, the executive director or another senior leader responded.

The survey was designed to collect agency perspectives on various topics, including how agencies are using, or considering using, AI to support their services and operations; challenges encountered; approaches to risk mitigation; and opportunities for additional state guidance and support.

JLARC received completed responses from all 80 agencies surveyed (100 percent response rate). In several cases where agencies shared IT infrastructure or leadership, such as the Virginia Economic Development Partnership and the Virginia Tourism Corporation, only one response was submitted.

Data collection and analysis

Analyses of VITA cloud access security broker data (Chapter 3)

To understand state employee use of general-purpose AI tools, JLARC received point-in-time summary AI usage data for state employees from VITA's cloud access security broker. JLARC received two summary datasets covering two time periods: April 12, 2026, to May 12, 2026, and July 27, 2026, to August 27, 2026. This data was available on a statewide basis only for VITA-supported agencies and contained several usage metrics, including the number of users and volume of data exchanged for each AI service. No individual-level or agency-level data was available.

JLARC staff compared usage of the VITA-provided general-purpose AI tool, Copilot Chat, with aggregate usage of other, publicly available general-purpose AI tools. Usage metrics for enterprise Copilot Chat would also capture any separate, paid Copilot Chat account usage but not usage of the free, publicly available Copilot Chat for Consumers. The aggregated usage metrics include ChatGPT, Claude, Gemini, and the free, publicly available Copilot Chat for Consumers.

Analyses of federal AI use case data (Chapter 4)

As part of its review of opportunities for AI use in Virginia, JLARC staff analyzed the U.S. Office of Management and Budget's (OMB) 2025 Federal Agency Artificial Intelligence Use Case Inventory, accessed August 17, 2026. Some federal agencies and use cases, such as those related to national security or basic research, are not subject to OMB's reporting requirement. The federal AI use case inventory has two components: individually reported use cases, which are agency-submitted profiles of specific AI systems, and consolidated reporting on federal agencies' use of common, commercial off-the-shelf AI products for a defined list of tasks.

Using the individually reported data, JLARC staff analyzed the number of use cases agencies reported as piloted or deployed by year, and how those use cases in 2025 were distributed across topic areas. Using the consolidated data, JLARC staff analyzed how many of the 45 reporting agencies used AI across each of OMB's 20 standard task categories. Both datasets are self-reported by agencies, and categorization choices (e.g., an agency's stage-of-development or topic-area designation) are agency-reported rather than independently verified.

Review of relevant national and international publications

JLARC staff reviewed research literature and publications on AI to identify key risks and benefits of AI use, relevant risk mitigation policies and processes, and factors affecting the cost and availability of AI systems. Sources included publications from:

- federal government entities, including the National Institute of Standards and Technology, the U.S. Office of Management and Budget, the Congressional Research Service, and the U.S. Government Accountability Office;
- international organizations, including the Organisation for Economic Co-operation and Development and the United Kingdom’s AI Security Institute;
- university-based research centers, including the Massachusetts Institute of Technology, Stanford University’s Institute for Human-Centered Artificial Intelligence and Cyber Policy Center, and Georgetown University’s Center for Security and Emerging Technology;
- research and policy organizations, including the Ada Lovelace Institute, Center for AI and Digital Policy, Center for Democracy & Technology, Pew Research Center, and RAND Corporation; and
- associations and technical organizations representing government and industry practitioners, including the National Association of State Chief Information Officers and the Open Worldwide Application Security Project.

Review of relevant laws, policies, and standards in Virginia and other states

JLARC staff also reviewed laws, policies, standards, and guidance governing AI use in Virginia and other states, including:

- Virginia laws relevant to state agency AI use and VITA’s statutory responsibilities;
- VITA’s AI policies, enterprise architecture standards, and guidance to agencies on adopting and using AI tools;
- other states’ policies and processes for guiding AI adoption by their state agencies;
- other states’ enacted and proposed legislation regulating AI use within and outside state government; and
- International Organization for Standardization and International Electrotechnical Commission (ISO/IEC) AI standards.

Use of AI tools

To support its research, JLARC staff used AI tools consistent with JLARC’s AI use policy. Agency general-purpose AI tools were used, for example, to help identify potential sources or research methods to consider and to make sentences clearer and more concise. All findings and recommendations reflect the independent judgment and analysis of JLARC staff.

Appendix C: Agency responses

As part of an extensive validation process, the state agencies and other entities that are subject to a JLARC assessment are given the opportunity to comment on an exposure draft of the report. JLARC staff sent an exposure draft of relevant portions to the secretary of administration, the Virginia Information Technologies Agency (VITA), the Virginia Department of Social Services, the State Corporation Commission, the Division of Legislative Services, and the Joint Commission on Technology and Science.

Appropriate corrections resulting from technical and substantive comments are incorporated in this version of the report. This appendix includes a response letter from the secretary of administration and VITA.



COMMONWEALTH of VIRGINIA

Office of the Governor

Traci J. Deshazor
Secretary of Administration

October 2, 2026

Mr. Hal E. Greer, Director
Joint Legislative Audit and Review Commission
919 E. Main Street, Suite 2101
Richmond, VA 23219

Dear Director Greer:

Thank you for the opportunity to comment on the exposure draft of JLARC's report, *Review of Artificial Intelligence in Virginia*. This administration is committed to a strong partnership with JLARC, the Joint Commission on Technology and Science, and the General Assembly, so that Virginia government can embrace technology innovation while managing the risks that come with it. The draft report provides valuable insights and recommendations for transformation and meaningful collaboration across government to address the significant risks presented by unregulated AI development, while also making AI work for Virginians' benefit.

AI is a horizontal technology development that has brought rapid and transformative change to every person, business, and government actor around the world. The unprecedented scale, impact, and adoption of AI demands a comprehensive response from a vast array of impacted stakeholders. The federal government should be the coordinating actor for AI regulation, but it has declined to act and has prevented state actions through threats to federal funding streams. In the absence of federal leadership, Governor Spanberger created the rapid-response AI Task Force—which I co-lead with the Chief Transformation Officer and Counsel to the Governor—to help safeguard people and communities. Public safety must always come first when developing technology, and AI is no exception.

As Secretary of Administration, I supervise the state agencies responsible for the core administrative and support functions that keep government running, including the Virginia Information Technologies Agency ("VITA"). But AI preparedness extends far beyond Virginia's government and VITA. Virginia's public safety agencies have a broad and critical role across both the private and public sectors, which is why Virginia's Unified Readiness Framework—established in Executive Order No. 19 (2026) and overseen by the Secretary of Public Safety and Homeland Security—expressly addresses critical infrastructure protection, cybersecurity, technology modernization, and data integration and information sharing. The AI Task Force is focused on bringing together all critical components across state government as well as the private sector to guard against AI risks, harden our infrastructure, and grow the Commonwealth's AI competency.

To be sure, VITA is a required and central participant in any discussion on AI cybersecurity and risk as it relates to Virginia's public sector. And the agency should be commended for taking the first steps in developing a coordinated AI

policy approach under the prior administration's executive order. But the issues presented by AI are too vast for any one agency to address and the existing approach to AI has failed to create internal processes and procedures that account for



COMMONWEALTH of VIRGINIA

Office of the Governor

Traci J. Deshazor
Secretary of Administration

the diverse perspectives, interests, and needs across state government. Those issues are clearly reflected in the state agency feedback included in the draft report, and I am committed to implementing JLARC's recommendations to significantly reform VITA's internal processes and client service. This administration intends for VITA to set the standard in AI risk mitigation for state systems while providing exceptional, flexible service and support that meets state agency needs promptly with minimum delay and unnecessary process.

As VITA assesses and implements the recommendations—such as the valuable recommendation to streamline internal processes and identify “root causes” of current impediments—it is important to also identify those causes that are beyond VITA's independent authority or resources to address. These likely include, for instance, statutory provisions, regulatory requirements, and resourcing considerations both at the state and individual agency level. These may relate to division of authorities across agencies, procurement standards, data and technology governance standards, appropriate division of cyber capabilities, human resource policy and hiring approaches across agencies, and resource constraints related state and federal appropriations. The full scope of relevant considerations is beyond any one agency's capacity to address.

In short, good governance should be efficient, citizen centered, safe, and reliable—our approach to AI is no different. As Secretary of Administration, I remain focused on making Virginia government work better for the people it serves. VITA will continue to play a critical role in AI safety, security, and innovation in state government, but the Commonwealth's holistic approach to AI extends far beyond a single agency. Through the AI Task Force—and leveraging the AI policy planning unit—this administration will review and address your recommendations, while determining how the necessarily multi-Secretariat and multi-agency approach to rapid AI development can best be implemented. JLARC's assessment will be a valuable addition to that comprehensive approach, and I appreciate the opportunity to review the draft report in advance.

Sincerely,

A handwritten signature in cursive script, reading "Traci J. Deshazor".

Traci J. Deshazor
Secretary of Administration
Commonwealth of Virginia



COMMONWEALTH of VIRGINIA

Virginia Information Technologies Agency

Michael Watson
Chief Information Officer
Email: cio@vita.virginia.gov

7325 Beaufont Springs Drive
Richmond, Virginia 23225
(804) 510-7300

TDD VOICE -TEL. NO.
711

October 2, 2026

Mr. Hal E. Greer, Director
Joint Legislative Audit and Review Commission
919 E. Main Street, Suite 2101
Richmond, VA 23219

Dear Director Greer:

Thank you for the opportunity to review and comment on the exposure draft of the forthcoming Joint Legislative Audit and Review Commission (JLARC) report, *Review of Artificial Intelligence in Virginia*. We appreciate the thoroughness of your work and agree with the overall recommendations as they relate to VITA's role in fostering innovation and enabling the safe and effective use of Artificial Intelligence (AI) across state agencies to make government work better for Virginians.

Since becoming the Chief Information Officer in April, I have made AI one of VITA's strategic priorities. To support this initiative VITA has created a Commonwealth of Virginia AI Strategic Framework to both operationalize the use of AI across executive branch agencies and improve the governance structure of the environment. While additional work remains to fully implement these recommendations, I would like to briefly highlight efforts VITA has already undertaken in recent months to create an AI-ready foundation for state agencies to adopt AI solutions. For example, as referenced in Recommendation 7, VITA is improving how it communicates to agencies the roadmap of AI tools and services. This effort will help agencies understand which AI tools are approved for state employee use as well as how to begin preparing for business use. Likewise, as referenced in Recommendation 8, VITA has recognized the need to train the workforce and proactively developed an employee AI training plan that will launch shortly.

Implementation of Recommendations 9 and 11 is also underway. Specifically, VITA is establishing an expedited review and approval determination process for AI coding assistance tools proposed by state agencies for IT system modernization. Additionally, VITA is working on developing a centralized, one-stop source for AI-related information, including tools under evaluation, tools in the POC stage, implemented use cases, and other relevant AI initiatives across the Commonwealth.

VITA strives to continuously improve on the efforts currently underway, and we agree with JLARC's assessment of what still needs to be accomplished. As such, we look forward to also

AN EQUAL OPPORTUNITY EMPLOYER


fully implementing Recommendations 1-6. VITA needs to strengthen our existing standards and use policies to ensure a consistent understanding of their meaning and application. We also need to create more streamlined, coordinated, and transparent processes that empower state agencies to better evaluate risk and more easily adopt AI solutions that safely meet business needs. VITA must establish clear processes to communicate AI progress and activities to our customers.

VITA embraces this role of establishing an AI-ready environment that will provide the Commonwealth's executive branch agencies with a solid foundation to adopt AI solutions. Fostering innovation and becoming AI-ready requires partnership with stakeholders, the General Assembly, the Joint Commission on Technology and Science, and JLARC.

Meeting the challenges and opportunities presented by AI and ensuring Virginia can be a national leader in AI preparedness requires a whole-of-government response and coordination among agencies and the Governor's Office. We look forward to engaging, advising and supporting the Secretariat of Administration, the Governor's Office, and the new AI Task Force, created by Executive Order No. 22 (2026), as they consider statewide AI policies and work to position the state to best meet the moment of AI.

We welcome additional feedback and collaboration as we work to implement the recommendations outlined in this report and continuously improve the services we provide our customer agencies to make government work better for Virginians.

Sincerely,

A handwritten signature in black ink that reads "Michael Watson". The script is cursive and fluid, with the first letters of each word being capitalized and prominent.

Michael Watson
Chief Information Officer

Appendix D: Key factors affecting cost and availability of AI systems

The resolution for this study directed JLARC to review how various factors could impact the cost and availability of AI. AI has a global supply chain, and many factors affect its cost and availability. Some key factors are:

- the highly concentrated supply chain for the hardware that AI runs on, particularly computer chips;
- constraints on the development of data centers, which house the hardware needed to power AI;
- the software and hardware efficiencies that determine how much hardware may be needed to complete a task; and
- the increasing availability and capabilities of open-weight models.

The cost and availability of an AI tool depend heavily on its scale and purpose. For example, a large organization using an AI tool for a high-risk, high-volume core business activity may require significant computational capacity and controls to ensure reliability and regulatory compliance. In contrast, an individual using an open-weight model for a personal offline summarization application will incur some energy and hardware costs but will not rely on a data center for that use. Given the complexity of these factors, it is difficult to predict how they may affect the costs and availability of AI in the future. This appendix summarizes published research and government reports on these factors.

Computer chip supply chain is highly concentrated

Computer chips are a key part of the hardware needed to train and run AI systems, and hardware availability and costs can have downstream effects on the costs incurred by businesses and individuals for AI systems and services. This section highlights key concentrations in the computer chip supply chain, though others exist.

There are two main types of chips used to train and run AI systems: processor chips, which perform the computations needed to generate AI outputs, and memory chips, which store and supply data to the processor chips. Currently, graphics processing units (GPUs) are the most commonly used processor chips for running AI systems, and the design and manufacturing of GPUs is highly concentrated. Just two companies are largely responsible for the design (NVIDIA) and production (Taiwan Semiconductor Manufacturing Company [TSMC]) of GPUs needed for AI compute, according to the Organisation for Economic Co-operation and Development (OECD) and the Stanford Institute for Human-Centered Artificial Intelligence (HAI). In addition to GPUs, leading AI developers are increasingly investing in more specialized processor chips known as application-specific integrated circuits, such as Google's tensor processing units and Amazon's Trainium and Inferentia chips.

High bandwidth memory (HBM) is used to store and supply data to the GPU and also has a highly concentrated supply chain. According to a 2025 OECD policy paper on competition in AI infrastructure, three companies (SK Hynix, Samsung, and Micron) collectively hold over 60 percent of the

market share for HBM. Because HBM is used in GPUs, any bottlenecks in HBM production can slow the supply chain for AI processor chips.

Certain critical minerals and materials needed for AI hardware, such as copper, silicon, gallium, and aluminum, are also concentrated geographically. According to the International Energy Agency's (IEA) 2025 Energy and AI report, three countries produce 60 percent of the world's refined copper, and China refines over 80 percent of the world's silicon, around 90 percent of the world's rare earth elements, and around 99 percent of the world's gallium.

Data center developments face increasing constraints

Data centers house the hardware, software, and networks that enable many AI systems to run. According to the IEA, AI model training and deployment occurred primarily in data centers as of 2025. However, data center developments are facing increasing constraints, particularly related to energy availability and community acceptance.

Data centers supporting AI systems can consume significant amounts of energy. "A typical AI-focused data center consumes as much electricity as 100,000 households," according to the IEA's 2025 Energy and AI report. Some planned data center campuses are expected to require more power than Virginia's largest nuclear reactor can generate, according to JLARC's 2024 report on data centers. Although the training of AI models consumes substantial energy, an increasing share of AI energy consumption comes from inference, the process of running a trained AI model after deployment. Within months, an AI system's inference energy demands can exceed its training energy demands, according to Stanford University's HAI 2026 AI Index report.

In some markets, energy demand for new data centers outstrips utilities' ability to provide the power requested. According to the IEA's 2025 Energy and AI report, connections are taking up to 10 years in some markets, and an estimated 20 percent of planned global data center buildout faces potential delays from energy grid bottlenecks. Meeting data centers' energy demand will require substantial amounts of new power transmission and generation infrastructure, but building this infrastructure at the pace data centers are demanding will be very difficult. (See JLARC's 2024 report, *Data Centers in Virginia*.)

Another constraint on data center development is community opposition. Communities are increasingly opposing data center projects over a variety of concerns, such as their impacts on the environment, electricity costs, and local residents. Between January 2026 and July 2026, legislation to place a moratorium on data center development had been introduced in 15 states, including Virginia, according to the National Conference of State Legislatures, although none had been enacted. Some Virginia localities have also moved to exert more control over data center development in 2026 through their zoning processes, including:

- Chesterfield County removed "by-right" data center development, requiring a conditional use approval through its zoning process;
- the City of Chesapeake recently enacted an eight-month pause on reviewing proposed data center applications and removed "by-right" data center development from its zoning code; and

- the City of Virginia Beach enacted a year-long pause on reviewing new large-scale and hyperscale data center applications; this pause may be shortened if data center regulations are updated in the city’s zoning code.

Data center water use has been a particular concern for some residents and environmental groups and can affect the location of data centers. Data centers require industrial-scale cooling, which is sometimes dependent on water, to manage the heat generated by their computing equipment. Whether data centers use significant amounts of water depends on the cooling technology used. In addition, whether significant water use is problematic depends on the local availability of water. For example, a recent Virginia Department of Environmental Quality study found that “under current conditions, it appears unlikely that a data center with evaporative cooling technology (or any comparable water user) would find a reliable, sufficient groundwater supply anywhere in the [Virginia Coastal Plain],” an aquifer system east of Interstate 95. However, a data center using different cooling technology or located in a region with sufficient groundwater may not face such constraints.

For more information about energy needs, water use, and other impacts from data centers in Virginia, see JLARC’s 2024 report, *Data Centers in Virginia*.

AI software and hardware efficiency gains have lowered AI computation costs for a given level of performance

AI software and hardware have become more efficient over time, reducing the cost of reaching a certain performance level. According to the 2026 International AI Safety Report, algorithmic efficiency improvements have historically reduced the compute needed to train a model to a specified performance level by a factor of two to six annually. Additionally, Epoch AI estimates that AI hardware has become 30 percent less expensive each year for a given level of hardware performance. A March 2026 study by MIT experts found that the price for a certain performance benchmark has decreased five to 10 times per year, which they attributed to “economic forces, hardware efficiency improvements, and algorithmic efficiency improvements.”

However, these efficiency gains may not result in reduced overall AI costs. For example, training costs for frontier AI models have increased 3.5 times annually since 2020, according to Epoch AI. Some experts suggest that efficiency gains may increase consumption rather than decrease it, ultimately increasing demand overall for hardware and data centers.

Growth in availability and capabilities of open-weight models is increasing access to AI to some extent

AI developers vary in the levels of access they permit to an AI system’s foundational components, including the numerical weights associated with the model’s neural network, the underlying code, and the training data. Access to the AI system components falls along a spectrum between “closed-source” and “open-source,” depending on how much of these fundamental components are made available to the public.

“Closed-source” models keep underlying components private; users cannot download or modify them. Examples of closed-source models are Anthropic’s Claude Opus 4.6, Google’s Gemini 2.5 Pro,

and OpenAI's GPT-5. In contrast, "open-source" models allow the public to access and download all fundamental components of the model, including the weights, code, and training data.

"Open-weight" AI models fall between "closed-source" and "open-source" models, making some of the foundational components free and available to download and modify, though some restrictions may apply. This allows users outside the model developer to customize and build on an AI model themselves, rather than interacting with the model through a controlled interface. With open-weight models, the developer releases the model's weights for download and modification, but not its code or training data.

Available open-weight models vary widely in size, so some require more memory, storage, and computing power than is typically available on personal computers. Examples of open-weight models available for the public to download and run on their personal computers include Google's Gemma 4 E2B and E4B (U.S. models), OpenAI's gpt-oss-20b (U.S. model), and Alibaba's Qwen3 8B (Chinese model).

In recent years, open-weight models have become increasingly available and capable. According to Stanford University's HAI 2026 AI Index report, "while closed-weight models still lead, open-weight models are far more competitive than they were a few years ago." For example, the report cited one common AI model performance benchmark in which users evaluate model responses showing the top performing open-weight model scored slightly higher in March 2026 than the top performing closed-weight model a year prior (Figure D-1). A more recent analysis by Epoch AI found that, since January 2026, the leading open-weight models have trailed the leading closed models by four months, on average, in overall capability.


Although performance varies depending on task and model, these trends indicate that AI capabilities are becoming more broadly available and accessible to the general public.

FIGURE D-1

Performance of open-weight models on one commonly cited benchmark is increasing

Performance of top closed vs. open models on the Arena

Source: Arena, 2026 | Chart: 2026 AI Index report



SOURCE: Stanford Institute for Human-Centered Artificial Intelligence (HAI), 2026 AI Index Report.

NOTE: Stanford's HAI cautions that while benchmarks remain useful indicators of technical progress, they are vulnerable to reliability problems, flawed questions, saturation, and gaming.

Appendix E: Recent AI legislation passed by the Virginia General Assembly

This appendix lists the AI-related legislation passed by the Virginia General Assembly during the 2024, 2025, and 2026 sessions (Table E-1). The table includes bills passed by both chambers but subsequently vetoed, and one that passed with a re-enactment clause.

TABLE E-1
AI-related bills passed between 2024 and 2026

Year	Bill Number (Patron)	Description
2024	SB 487 (Aird)	<i>Study</i> - Directs the Joint Commission on Technology and Science (JCOTS) to study <u>artificial intelligence use by public bodies in Virginia</u> .
2024	SB 731 (Durant)	Amends the existing child sexual abuse material statute to clarify that the child depicted does not have to exist, criminalizing <u>AI-generated child sexual abuse material</u> .
2024	Appropriation Act (Item 426 N)	Authorizes the Department of Motor Vehicles commissioner "to issue a Request for Information for (i) updating customer-facing web applications; (ii) pursuing the use of <u>artificial intelligence in day-to-day activities</u> ; (iii) the issuance of digital passports and mobile driver's licenses; (iv) improving customer service, specifically through smart phone technologies and the use of self-service kiosks; and (v) other innovative technologies to improve the overall customer experience."
2025	HB 1642 (Hayes)	Requires all decisions related to the pre-trial detention or release, prosecution, adjudication, sentencing, probation, parole, correctional supervision, or rehabilitation of criminal offenders to be made by the judicial officer or other person charged with making such decision. Bill specifies that these decisions may not be made without involvement of a human decision-maker and that the use of any recommendation or prediction from an artificial intelligence-based tool is subject to any challenge or objection permitted by law.
2025	HB 2692 (Glass)	Prohibits law-enforcement officers from knowingly and intentionally making false statements about any known material fact, including by use of inauthentic replica documents ("any documents, including <u>computer-generated documents, created by any means, including artificial intelligence</u> , by a law-enforcement officer or his agent that (i) contain a false statement, signature, seal, letterhead, or contact information or (ii) materially misrepresent any fact") prior to or during a custodial interrogation of a child to secure the cooperation, confession, or conviction of such child.
2025	HB 2124 (Maldonado)/ SB 1053 (Ebbin)	Applies provisions related to defamation, slander, and libel to synthetic digital content (any AI-produced digital content that "realistically but falsely depicts an individual's appearance, speech, or conduct"). <i>Note: This bill passed with a re-enactment clause.</i>

2025	HB 2094 (Maldonado)	<u>Establishes operating standards for developers and employers of high-risk AI systems</u> (AI systems “specifically intended to autonomously make, or be a substantial factor in making, a consequential decision”), including implementing risk management programs, conducting impact assessments, providing transparency documentation regarding performance and known risks for algorithmic discrimination, and notifying individuals when a high-risk AI system is used to make a consequential decision concerning them. <i>Note: This bill passed both chambers but was vetoed by the governor.</i>
2025	HB 2479 (Sickles)/ SB 775 (Surovell)	Prohibits political ads containing synthetic media, such as AI-generated deepfakes, unless they include a clear disclaimer about the use of synthetic media. <i>Note: This bill passed both chambers but was vetoed by the governor.</i>
2026	HB 171 (Askew)	Updates Virginia’s school internet-safety law to require local school boards’ student internet-safety instruction to cover modern digital safety topics, including online scams, misinformation, and content generated by <u>artificial intelligence</u> . It also directs the superintendent of public instruction to issue guidance to school divisions on those topics.
2026	HB 580 (Glass)	Requires the Office of the Attorney General to “establish mechanisms for receiving and investigating complaints by the Commonwealth’s consumers <u>involving emerging technologies</u> , including referring appropriate complaints to the federal, state, and local departments or agencies charged with the enforcement of applicable consumer laws.”
2026	HB 797 (Hayes)/SB 384 (Graves)	<i>Study</i> - Directs JCOTS “to evaluate and report on the feasibility and impact of developing a framework for any person or entity seeking to act as an <u>independent verification organization that assesses artificial intelligence models’ or applications’ adherence to standards</u> reflecting best practices for the prevention of personal injury and property damage.” “If JCOTS recommends the development of such a framework, JCOTS shall further evaluate (a) the feasibility of a licensing scheme for independent verification organizations and (b) the viability of the Virginia Information Technologies Agency or any other state agency to implement and manage such licensing scheme.”
2026	HB 1186 (Rasoul)/SB 394 (Pekarsky)	Requires the Virginia Department of Education (VDOE) to (i) gather information on how AI systems are used for student instruction in Virginia’s public schools and (ii) create and share guidance on the “safe, ethical, and equitable use” of AI systems (AIS) used for student instruction in public elementary and secondary schools. Subject to appropriations, directs VDOE “to establish and oversee the AIS Innovation in Education Pilot Program for the purpose of funding, evaluating, and scaling innovative uses of AIS in public elementary and secondary schools by providing support to school divisions in piloting AIS applications for instruction, tutoring, student engagement, operational efficiency, and teacher support and to submit an annual report to the Chairs of the House Committee on Education and the Senate Committee on Education and Health by December 1.”
2026	HJ 32 (Laufer)	<i>Study</i> - Requires the Joint Legislative Audit and Review Commission to study current <u>artificial intelligence</u> use <u>policies</u> at higher education institutions in Virginia.

2026	Appropriation Act (Item 113 S.6)	Allocates \$500K in FY27 to the MOVE Chamber, in collaboration with the Virginia Economic Development Partnership and the Virginia Community College System, “to develop and implement a pilot program to provide internship and apprenticeship opportunities with existing Virginia businesses for at least 1,000 students and recent graduates of the Virginia Community College System in Artificial Intelligence, Data Science, and Cybersecurity over the next three years.”
2026	Appropriation Act (Item 295 T.4)	Directs the Department of Medical Assistance Services to “seek opportunities to enhance call center operations through the use of <u>artificial intelligence (AI)</u> . All call center related contract procurements, re-procurements, or modifications shall maximize the use of AI to reduce costs and improve service. Any use of AI must ensure the protection of personal information and comply with federal law and regulations.”
2026	Appropriation Act (Item 1 N.5)	Provides the Joint Subcommittee on Tax Policy with several goals and objectives for the 2026 interim, including to “study the data center sales and use tax exemption” and to <u>evaluate AI’s impact on employment</u> .
2026	Appropriation Act (Item 148 N)	Allocates \$500K in FY27 for George Mason University to develop and scale the AI Innovation Nexus.
2026	Appropriation Act (Item 188 F)	Allocates \$500K in FY27 for University of Virginia “to coordinate the planning and development of the Commonwealth AI Institute.”
2026	Appropriation Act (Item 124 HHH)	Allocates \$2M in FY27 for VDOE “to support the AI Innovation in Education Pilot Program. The Department shall administer the Program and develop guidelines. Funding may be used to support the adoption of artificial intelligence software, programs, and training.”

SOURCE: JLARC staff analysis of 2024–26 legislation in Virginia.

Appendix F: Risk-based approaches to AI oversight in other states

This appendix provides examples of how other states differentiate approval and oversight requirements for state agency use of AI based on risk.

Maryland

As part of Maryland's AI intake process, which is outlined in its Responsible AI Policy and implementation guidance, agencies are responsible for initially classifying a proposed AI project into one of four risk tiers based on its data sensitivity, impact on constituents, and the AI's criticality:

- **Minimal risk**: AI systems that pose a negligible risk, are embedded in standard office software or are used for purely internal tasks with no impact on individual safety or rights, and process public or "Protected/Internal Use Only" data.
- **Limited risk**: AI systems that have moderate or low impact, few inherent risks, and process public or "Protected/Internal Use Only" data.
- **High risk**: AI systems that may significantly affect individuals or critical government operations or that process confidential or restricted data.
- **Unacceptable risk**: AI systems that pose extreme risks to public welfare, safety, or rights that cannot be adequately mitigated and are banned from use.

Systems designated as high risk receive in-depth review and must demonstrate they have robust safeguards. Agencies must complete algorithmic impact and privacy impact assessments, implement risk mitigation measures, perform ongoing monitoring, and maintain alignment with AI risk management best practices from the National Institute of Standards and Technology. Prior to full deployment, high-risk uses may be required to first conduct a proof of concept or add additional safeguards. Maryland law also requires agencies using high-risk AI to conduct regular impact assessments, which the guidance states can be done by re-evaluating the system at least annually using the algorithmic impact assessment.

In contrast, limited-risk systems require only "light-touch" safeguards, including transparency requirements (e.g., disclosing use in chatbots) and basic testing for fairness and accuracy, while minimal-risk uses require no AI-specific approvals beyond standard security and IT reviews.

Idaho

Idaho's governance model for state agency AI use is centered on risk-based oversight, where all AI systems are evaluated for risk based on six factors: data sensitivity, decision impact, autonomy level, transparency, scope and scale, and novelty and complexity. Each system is rated on a four-point scale for each factor, producing a composite risk score that aligns with one of the three governance tiers:

- **Tier 1 (Low risk)**: AI systems that generally process non-sensitive information, have limited decision impact, maintain human oversight, and have a narrow scope of operation.

- **Tier 2 (Medium risk):** AI systems that may process some sensitive information, have moderate decision impact, operate with reduced human oversight, or serve a broader user base.
- **Tier 3 (High risk):** AI systems that process highly sensitive information, have significant decision impact, operate with substantial autonomy, or serve large populations.

AI systems are subject to different levels of approval based on their risk tier. For example, Tier 1 systems require approval by only the agency's IT leader and the central IT agency, while Tier 3 systems require approval from agency leadership and the central IT agency, as well as mandatory consultation from the Ethics Committee, Technical Review Board, and Executive Committee.

Beyond approval procedures, AI systems are also subject to a variety of requirements based on their assessed risk level, including requirements related to documentation, security, and monitoring (Table F-1).

TABLE F-1
Examples of Idaho's risk-tiered requirements

Requirement	Tier 1	Tier 2	Tier 3
Register in AI system inventory	X	X	X
AI concept brief documenting purpose, anticipated benefits, and alignment with agency mission	X	X	X
Implement access controls for system administration and operation	X	X	X
Maintain system logs	X	X	X
Disclose when generative AI contributes to public-facing content	X	X	X
Privacy impact assessment ^a addressing data collection, use, storage, and sharing		X	X
Document fairness evaluation results across demographic groups		X	X
Enhance access controls with stronger authentication		X	X
Detailed model documentation including architecture, training methodology, and limitations			X
Document risk mitigation plan			X
Enhanced security assessment documenting specialized AI security controls			X
Conduct regular penetration testing			X
Performance review	Quarterly	Monthly	Weekly
Risk reassessment	Annual	Semi-annual	Quarterly

SOURCE: JLARC summary of Idaho's *Artificial Intelligence Governance Policy, Standard, and Guideline* (2025)

NOTE: This is not a comprehensive list of all of Idaho's risk-tiered requirements. ^a Privacy impact assessments are required for any AI system processing personal information.

Texas

Texas law requires identification of “heightened scrutiny AI systems,” defined as AI systems intended to autonomously make a consequential decision, serve as the principal basis for such a decision, or alter the outcome of such a decision. The definition excludes systems intended to perform a narrow procedural task; improve the result of a previously completed human activity; perform a preparatory task to an assessment that informs a consequential decision; or detect decision-making patterns or deviations from previous decision-making patterns.

In addition to requirements imposed on all AI systems used by state agencies (e.g., systems must allow humans to review inputs and outputs, disclosure requirements for public-facing systems), Texas imposes additional, more rigorous requirements for heightened scrutiny AI systems. Before deploying a heightened scrutiny AI system or if a material change is made to the system, agencies must complete a written AI risk assessment that considers the probability and severity of harm that could occur. The assessment documents known security risks and mitigation steps; accuracy and efficiency metrics; decision-making methodology and training data; and the availability of inputs and outputs to monitor the system’s decision-making over time. Based on this assessment and the identified risk mitigation measures, the agency’s designated AI risk officer must approve deployment of the heightened scrutiny AI system.

Agencies must also complete impact assessments for heightened scrutiny AI systems prior to their deployment or when a material change is made to the system, the state or local data used by the system, or its intended use. The impact assessment requires documenting the system’s training data, model, and intended use; the agency staff responsible for deployment and ongoing monitoring; whether the system will process or store any personally identifiable information; potential risks of unlawful harm and steps to limit those risks; identified system limitations; monitoring procedures and intervals; and data retention procedures.

Agencies must also identify the acceptable use cases for a heightened scrutiny system and adopt an acceptable use policy that prevents uses other than those approved by the agency for that system. All employees must be adequately trained on the acceptable use policy, and employees or contractors who access, use, or manage the system must be provided training on identified risks and appropriate methods for managing those risks.



JLARC.VIRGINIA.GOV

919 East Main Street Suite 2101 Richmond, VA 23219