

HOUSE BILL NO. 797
AMENDMENT IN THE NATURE OF A SUBSTITUTE
(Proposed by the House Committee on Appropriations
on _____)
(Patron Prior to Substitute—Delegate Hayes)

A BILL to amend the Code of Virginia by adding a section numbered 2.2-2012.01 and by adding in Chapter 20.1 of Title 2.2 an article numbered 9, consisting of sections numbered 2.2-2034.2 through 2.2-2034.8, relating to Virginia Information Technologies Agency; artificial intelligence; independent verification organizations.

Be it enacted by the General Assembly of Virginia:

1. That the Code of Virginia is amended by adding a section numbered 2.2-2012.01 and by adding in Chapter 20.1 of Title 2.2 an article numbered 9, consisting of sections numbered 2.2-2034.2 through 2.2-2034.8, as follows:

§ 2.2-2012.01. *Additional duties of the CIO relating to licensing of independent verification organizations of artificial intelligence.*

A. The CIO shall have the power and duty to oversee the licensing of independent verification organizations (IVOs) of artificial intelligence pursuant to Article 9 (§ 2.2-2034.2 et seq.). The CIO shall promulgate regulations necessary or incidental to the licensing of such IVOs, which shall include:

1. Conflict of interest and funding transparency requirements, including reporting requirements regarding the IVOs' funding sources and revenue generation and self-audit requirements regarding the IVOs' board composition to ensure adequate independence from the artificial intelligence industry;

2. Requirements for identifying additional IVO plan elements as needed to ensure acceptable levels of risk from IVO-verified artificial intelligence models or applications;

3. Provisions on circumstances mandating corrective action or loss of license;

4. Requirements related to the structure and terms of the Artificial Intelligence Safety Advisory Council, including the procedures for appointing additional members; and

5. Requirements related to IVO application procedures and required materials.

B. In promulgating regulations pursuant to this section and in administering Article 9 (§ 2.2-2034.2 et seq.) of Chapter 20.1, the CIO shall consult with the Consumer Protection Division of the Office of the Attorney General and the Artificial Intelligence Safety Advisory Council established pursuant to § 2.2-2034.7.

Independent Verification Organizations of Artificial Intelligence.

§ 2.2-2034.2. Definitions.

As used in this article, unless the context requires a different meaning:

"Artificial intelligence application" means a software program or system that uses artificial intelligence models to perform tasks that typically require human intelligence.

"Artificial intelligence model" means an engineered or machine-based system that can, for explicit or implicit objectives, infer from the input it receives how to generate outputs that can influence physical or virtual environments.

"Deployer" means a person or entity that implements, integrates, or makes operational an artificial intelligence model or artificial intelligence application within the Commonwealth. "Deployer" includes a person or entity that makes an artificial intelligence model or artificial intelligence application available for use by others within the Commonwealth, whether directly or as part of a product or service.

"Developer" means a person or entity that develops an artificial intelligence model or artificial intelligence application that is deployed in the Commonwealth.

"Independent verification organization" or "IVO" means a person or entity licensed by the Virginia Information Technologies Agency (VITA) pursuant to this article to assess artificial intelligence models' or applications' adherence to standards reflecting best practices for the prevention of personal injury and property damage.

"Security vendor" means a third-party entity engaged by an IVO or developer to evaluate the safety or security of an artificial intelligence model or application, including by using processes such as red teaming, risk detection, and risk mitigation.

§ 2.2-2034.3. Licensing of independent verification organizations.

A. VITA may, but is not required to issue a license to a person or entity seeking to act as an IVO within the Commonwealth.

B. Any person or entity seeking to act as an IVO within the Commonwealth and receive a license for such activities shall submit an application to VITA. An application for an IVO license shall be made by filing with VITA the information, materials, and forms required by this article and the regulations promulgated by the CIO pursuant to § 2.2-2012.01, along with a plan detailing all of the following information:

1. The risks with respect to which the applicant intends to verify that artificial intelligence models or artificial intelligence applications implement mitigation measures sufficient to achieve acceptable levels of risk. For each such risk, the applicant shall submit (i) a proposed definition of acceptable levels of risk; (ii)

metrics that are measurable and can be used to determine whether the acceptable level of risk defined by the IVO produces beneficial outcomes; (iii) target levels for the metrics, including data sources such levels are based on and methods for measurement; and (iv) a description of the evaluation and reporting protocol to determine whether verified models or applications meet the outcome metrics on an ongoing basis.

2. Proposed technical, operational, governance, and other mitigation requirements for developers or deployers, including procedures for pre-development and post-development, to ensure acceptable levels of risk, including ongoing monitoring of risks and assessment of mitigation efficacy.

3. Methodologies and sources used to evaluate the efficacy of mitigation requirements and updates to such methodologies and sources as needed.

4. Benchmarks, technologies, and audit methodologies proposed to assess developer and deployer adherence to mitigation requirements.

5. Approach to assessing continued good standing of a developer or deployer, including reviewing and evaluating the developer's or deployer's maintenance of artificial intelligence governance plans and policies, processes for risk monitoring and mitigation, whistleblower protections, and training for employees and third parties.

6. Disclosure requirements for developers or deployers related to detected risks, incident reports, or material changes to risk profiles, including both risks detected prior to verification and risks resulting from fine-tuning or modifying an artificial intelligence model or artificial intelligence application after verification.

7. Procedures for prescribing and verifying implementation of corrective actions to remedy an identified failure by a developer or deployer to do any of the following: (i) achieve an acceptable level of risk with respect to an artificial intelligence application or artificial intelligence model; (ii) comply with any other mitigation requirements promulgated by the applicant; or (iii) comply with the developer's or deployer's artificial intelligence governance plans and policy.

8. Standards and procedures for revoking verification for noncompliance with the applicant's mitigation requirements, failure to achieve acceptable levels of risk, or noncompliance with the developer's or deployer's artificial intelligence governance plans and policies.

9. Whether the applicant proposes market-specific verification and how plans are tailored to that segment.

10. Coordination with federal and state authorities.

11. Personnel qualifications.

12. Governance policies, sources of funding, and policies ensuring independence.

13. Any other information required by VITA.

C. VITA may license an applicant as an IVO if (i) such applicant demonstrates independence from the artificial intelligence industry and (ii) every element of the applicant's submitted plan is adequate to ensure that artificial intelligence models or artificial intelligence applications verified pursuant to such plan will mitigate to an acceptable level one or more risks including as defined by the metrics the applicant proposes. If verification is proposed by an applicant for a specific market segment, VITA shall evaluate the applicant's submitted plan accordingly. If VITA finds that an applicant's plan adequately mitigates some, but not all, of the proposed risks, the applicant shall be licensed to verify only those risks for which the plan is deemed adequate.

D. An IVO issued to an applicant by VITA shall specify the risks the IVO is authorized to verify and any market segments for which the license applies.

E. VITA shall establish reasonable application and renewal fees sufficient to offset administrative costs. Such fees shall be payable to VITA and used for (i) application processing, (ii) audits of IVOs, (iii) compensation of the Artificial Intelligence Safety Advisory Council, and (iv) general administration of this article.

§ 2.2-2034.4. License revocation.

A. VITA shall revoke an IVO license if it determines any of the following:

1. The IVO's plan is materially misleading or inaccurate;

2. The IVO fails to adhere to its plan in a way that materially impairs its responsibilities, including failure to adhere to the plan's procedures for ongoing monitoring of verified artificial intelligence models or applications and implementation of corrective action;

3. A material change compromises independence from the artificial intelligence industry;

4. Technological evolution renders methods obsolete for ensuring acceptable levels of the risk VITA has designated the independent verification organization to verify; or

5. A verified model or application causes material harm of the type the IVO defines an acceptable level of risk in order to prevent.

B. If VITA determines the public interest so requires, it may allow an IVO to cure the basis for revocation before terminating the license.

§ 2.2-2034.5. Independent verification organization responsibilities; modifications to plans.

A. A licensed IVO shall implement the approved plan submitted pursuant to subsection B of § 2.2-2034.3,

which includes verifying artificial intelligence models or artificial intelligence applications. An IVO shall revoke verification of an artificial intelligence model or artificial intelligence application if a developer or deployer (i) fails to meet mitigation requirements, (ii) fails to cooperate with monitoring, (iii) violates governance policies, or (iv) fails to implement corrective actions.

B. An IVO may (i) update or modify the following aspects of the approved plan: (a) technical and operational requirements; (b) evaluation benchmarks; (c) audit methodologies; (d) governance plans; or (e) any other element of its plan in order to take advantage of improved technology; (ii) address previously discovered issues with its plan; or (iii) otherwise enhance the efficacy of its verification activities. An IVO shall provide written notice to VITA of any material change to an approved plan, which shall describe the change, the rationale for the change, and the IVO's basis for concluding that the change will better ensure that verified artificial intelligence systems or artificial intelligence applications achieve acceptable levels of risk. For purposes of this subsection, "material change" means any change that alters mitigation requirements, acceptable levels of risk, benchmarks, audit methodologies, or governance plans in a manner that could reasonably affect the level of risk posed by verified artificial intelligence systems or artificial intelligence applications.

C. Within 180 days after receiving notice of a material change under subsection B, VITA may request additional information from the IVO or issue a written notice disapproving the material change in whole or in part if VITA determines that the change is inadequate to ensure acceptable levels of risk or is otherwise inconsistent with this article or regulations promulgated pursuant to § 2.2-2012.01. If VITA disapproves a material change, the IVO shall cease implementing the disapproved portion of the change and shall modify its plan accordingly.

D. If VITA disapproves a material change under subsection C or determines that a material change reveals that previously verified artificial intelligence systems or artificial intelligence applications may no longer achieve acceptable levels of risk, the IVO, in consultation with VITA, shall assess whether the previously verified artificial intelligence systems or artificial intelligence applications shall be re-assessed or subject to additional mitigation requirements and shall revoke or condition verification as appropriate.

E. Nothing in this article shall be construed to require any developer or deployer to seek verification by an independent verification organization for any artificial intelligence system or artificial intelligence application.

§ 2.2-2034.6. Annual reporting.

A. A licensed IVO shall submit an annual report to VITA including:

157 1. Aggregated information on the capabilities of the artificial intelligence models and artificial
158 intelligence applications evaluated by the IVO, the observed societal risks and benefits associated with those
159 capabilities, and the potential societal risks and benefits associated with those capabilities;

160 2. Adequacy of evaluation resources, technical capabilities, and mitigation measures to address observed
161 and potential risks;

162 3. Aggregated results of verification assessments;

163 4. Aggregated and anonymized compliance with prescribed remediation;

164 5. Anonymized descriptions of any additional, significant risk the IVO observed while conducting its
165 assessments, even if such risk is not one the IVO is licensed to verify;

166 6. A list of verified artificial intelligence systems;

167 7. A description of evaluation methods; and

168 8. Governance or funding changes affecting independence.

169 B. An IVO may redact trade secrets, sensitive business information, personally identifiable information,
170 and other security-sensitive content.

171 C. Documentation used in reports shall be retained for 10 years. A licensed IVO shall also retain all
172 documentation relating to its assessment and verification of artificial intelligence models or applications,
173 including ongoing monitoring and any subsequent corrective action, for 10 years following the relevant
174 activity.

175 D. VITA shall publish redacted versions of reports submitted by licensed IVOs pursuant to this section
176 online.

177 **§ 2.2-2034.7. Artificial Intelligence Safety Advisory Council; purpose, membership.**

178 A. The Artificial Intelligence Safety Advisory Council (the Advisory Council) is established as an advisory
179 council in the executive branch of state government for the purpose of advising and assisting VITA in
180 licensing and auditing IVOs. The CIO may delegate to the Advisory Council such powers and duties under
181 this article as the CIO deems appropriate, including participation in the evaluation of applications for
182 licensure, renewal, and revocation of licenses issued pursuant to §§ 2.2-2034.3 and 2.2-2034.4.

183 B. The Advisory Council shall have a total membership of up to 12 nonlegislative citizen members and
184 one ex officio member. Nonlegislative citizen members shall be appointed by the CIO and shall include at
185 least one citizen representative from a nongovernmental organization, educational and research institution,
186 public policy institute, or consumer and business advocacy organization. All members of the Advisory
187 Council shall be qualified to assess IVO plans. The Attorney General, or his designee, shall serve ex officio

with voting privileges and shall serve a term coincident with his term of office. Nonlegislative citizen members shall serve terms of four years and no member may serve more than two consecutive terms. Nonlegislative citizen members may be removed for inefficiency, neglect, or malfeasance. A majority of the members of the Advisory Council shall constitute a quorum. Nonlegislative citizen members of the Advisory Council shall be citizens of the Commonwealth and shall serve without compensation but shall be reimbursed for all reasonable and necessary expenses incurred in the performance of their duties as provided in §§ 2.2-2813 and 2.2-2825. Funding for the costs of expenses of the members shall be provided by VITA.

C. Nonlegislative citizen members of the Advisory Council shall (i) refrain from any action or occupation, gainful or not, that is incompatible with their duties, including employment by a developer or deployer of artificial intelligence; (ii) refrain from owning or acquiring any equity or other interest, directly or indirectly, in companies whose business consists in significant part in developing or deploying artificial intelligence; and (iii) observe a one-year post-employment restriction from employment by any artificial intelligence firms or IVOs.

D. The Advisory Council shall keep a record of its proceedings, including any considerations relating to the issuance, refusal, renewal, or revocation of IVO licensure.

§ 2.2-2034.8. Limitation of liability; indemnification.

A. No civil action shall lie against the Commonwealth, VITA, the CIO, the Artificial Intelligence Safety Advisory Council, or any member, officer, employee, or agent of any of the foregoing for any act or omission done in the performance of duties under this article, including any act or omission relating to the review, approval, denial, renewal, suspension, or revocation of a license issued pursuant to § 2.2-2034.3, the oversight or auditing of an independent verification organization, or the receipt, use, or publication of information or reports submitted pursuant to this article.

B. The Commonwealth, VITA, the CIO, the Artificial Intelligence Safety Advisory Council, or any member, officer, employee, or agent of any of the foregoing shall not be liable for any damages, injuries, or losses claimed to have resulted from any artificial intelligence system or artificial intelligence application that is verified or not verified by an independent verification organization, or from any decision to approve, deny, renew, suspend, or revoke a license of an independent verification organization.

C. The Commonwealth shall defend and indemnify any member of the Artificial Intelligence Safety Advisory Council and any officer, employee, or agent of VITA for any claim, demand, suit, or judgment arising out of any act or omission performed in good faith in the course and scope of duties under this article, to the same extent and subject to the same limitations as provided for state officers and employees under the

219 *laws of the Commonwealth.*

220 **2. That the Virginia Information Technologies Agency may accept any gifts, donations, grants,**
221 **bequests, or other private funds to effectuate the provisions of this act.**