

25105027D

HOUSE BILL NO. 2094

AMENDMENT IN THE NATURE OF A SUBSTITUTE

(Proposed by the House Committee on Communications, Technology and Innovation on January 27, 2025)

(Patron Prior to Substitute—Delegate Maldonado)

A BILL to amend the Code of Virginia by adding in Title 59.1 a chapter numbered 58, consisting of sections numbered 59.1-607 through 59.1-613, relating to high-risk artificial intelligence; development, deployment, and use; civil penalties.

Be it enacted by the General Assembly of Virginia:

1. That the Code of Virginia is amended by adding in Title 59.1 a chapter numbered 58, consisting of sections numbered 59.1-607 through 59.1-613, as follows:

CHAPTER 58.

HIGH-RISK ARTIFICIAL INTELLIGENCE DEVELOPER AND DEPLOYER ACT.

§ 59.1-607. Definitions.

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

"Algorithmic discrimination" means the use of an artificial intelligence system that 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, sexual orientation, veteran status, or other classification protected under state or federal law. "Algorithmic discrimination" does not include (i) the offer, license, or use of a high-risk artificial intelligence system by a developer or deployer for the sole purpose of the developer's or deployer's self-testing to identify, mitigate, or prevent discrimination or otherwise ensure compliance with state and federal law; (ii) the expansion of an applicant, customer, or participant pool to increase diversity or redress historical discrimination; or (iii) an act or omission by or on behalf of a private club or other establishment not in fact open to the public, as set forth in Title II of the Civil Rights Act of 1964, 42 U.S.C. § 2000a(e), as amended from time to time.

"Artificial intelligence system" means any machine learning-based system that, for any explicit or implicit objective, infers from the inputs such system receives how to generate outputs, including content, decisions, predictions, and recommendations, that can influence physical or virtual environments. "Artificial intelligence system" does not include any artificial intelligence model that is used for development, prototyping, and research activities before such artificial intelligence model is released on the market.

"Consequential decision" means any decision that has a material legal, or similarly significant, effect on the provision or denial to any consumer of (i) parole, probation, a pardon, or any other release from incarceration or court supervision; (ii) education enrollment or an education opportunity; (iii) access to employment; (iv) a financial or lending service; (v) access to health care services; (vi) housing; (vii) insurance; (viii) marital status; or (ix) a legal service.

"Consumer" means a natural person who is a resident of the Commonwealth and is acting only in an individual or household context. "Consumer" does not include a natural person acting in a commercial or employment context.

"Deployer" means any person doing business in the Commonwealth that deploys or uses a high-risk artificial intelligence system to make a consequential decision in the Commonwealth.

"Developer" means any person doing business in the Commonwealth that develops or intentionally and substantially modifies a high-risk artificial intelligence system that is offered, sold, leased, given, or otherwise provided to consumers in the Commonwealth.

"General-purpose artificial intelligence model" means a model used by an artificial intelligence system or other system that (i) displays significant generality, (ii) is capable of competently performing a wide range of distinct tasks, and (iii) can be integrated into a variety of downstream applications or systems.

"General-purpose artificial intelligence model" does not include any artificial intelligence model that is used for development, prototyping, and research activities before such artificial intelligence model is released on the market.

"Generative artificial intelligence" means an artificial intelligence system that is capable of producing and used to produce synthetic content, including audio, images, text, and videos.

"Generative artificial intelligence system" means any artificial intelligence system or service that incorporates generative artificial intelligence.

"High-risk artificial intelligence system" means any artificial intelligence system that is specifically intended to autonomously make, or be a substantial factor in making, a consequential decision. A system or service is not a "high-risk artificial intelligence system" if it is intended to (i) perform a narrow procedural task, (ii) improve the result of a previously completed human activity, (iii) detect any decision-making patterns or any deviations from pre-existing decision-making patterns, or (iv) perform a preparatory task to

HOUSE SUBSTITUTE

HB2094H1

1/27/25 11:36

60 *an assessment relevant to a consequential decision. "High-risk artificial intelligence system" does not include*
61 *any of the following technologies:*

62 *1. Anti-fraud technology that does not use facial recognition technology;*

63 *2. Anti-malware technology;*

64 *3. Anti-virus technology;*

65 *4. Artificial intelligence-enabled video games;*

66 *5. Calculators;*

67 *6. Cybersecurity technology;*

68 *7. Databases;*

69 *8. Data storage;*

70 *9. Firewall technology;*

71 *10. Internet domain registration;*

72 *11. Internet website loading;*

73 *12. Networking;*

74 *13. Spam and robocall filtering;*

75 *14. Spell-checking technology;*

76 *15. Spreadsheets;*

77 *16. Web caching;*

78 *17. Web hosting or any similar technology;*

79 *18. Autonomous vehicle technology; or*

80 *19. Technology that communicates with consumers in natural language for the purpose of providing users*
81 *with information, making referrals or recommendations, and answering questions and is subject to an*
82 *acceptable use policy that prohibits generating content that is discriminatory or unlawful.*

83 *"Integrator" means a natural or legal person that knowingly integrates a high-risk artificial intelligence*
84 *system into a software application and places said integration on the market under that person's mark,*
85 *whether for payment or free of charge. An integrator is neither a developer nor a deployer, except as*
86 *otherwise provided in § 59.1-610, nor will any person be deemed an integrator as a result of offering*
87 *information technology infrastructure.*

88 *"Intentional and substantial modification" means any deliberate change made to (i) an artificial*
89 *intelligence system that results, at the time when the change is implemented and any time thereafter, in any*
90 *new material risk of algorithmic discrimination or (ii) a general-purpose artificial intelligence model that*
91 *affects compliance of the general-purpose artificial intelligence model, materially changes the purpose of the*
92 *general-purpose artificial intelligence model, or results in any new reasonably foreseeable risk of*
93 *algorithmic discrimination. "Intentional and substantial modification" does not include (a) any customization*
94 *made by deployers based on legitimate nondiscriminatory business justifications and within the scope and*
95 *purpose of the artificial intelligence tool; (b) any change made to a high-risk artificial intelligence system, or*
96 *the performance of a high-risk artificial intelligence system, if (1) the high-risk artificial intelligence system*
97 *continues to learn after such high-risk artificial intelligence system is offered, sold, leased, licensed, given, or*
98 *otherwise made available to a deployer, or deployed, and (2) such change (A) is made to such high-risk*
99 *artificial intelligence system as a result of any learning described in clause (1) and (B) was predetermined by*
100 *the deployer or the third party contracted by the deployer and included within the initial impact assessment*
101 *of such high-risk artificial intelligence system as required in § 59.1-609.*

102 *"Machine learning" means the development of algorithms to build data-derived statistical models that are*
103 *capable of drawing inferences from previously unseen data without explicit human instruction.*

104 *"Person" includes any individual, corporation, partnership, association, cooperative, limited liability*
105 *company, trust, joint venture, or any other legal or commercial entity and any successor, representative,*
106 *agent, agency, or instrumentality thereof. "Person" does not include any government or political subdivision.*

107 *"Principal basis" means the use of an output of a high-risk artificial intelligence system to make a*
108 *decision without (i) human review, oversight, involvement, or intervention or (ii) meaningful consideration by*
109 *a human.*

110 *"Red-teaming" means adversarial testing to identify the potential adverse behaviors or outcomes of an*
111 *artificial intelligence system, identify how such behaviors or outcomes occur, and stress test the safeguards*
112 *against such behaviors or outcomes.*

113 *"Significant update" means any new version, new release, or other update to a high-risk artificial*
114 *intelligence system that results in significant changes to such high-risk artificial intelligence system's use*
115 *case or key functionality and that results in any new or reasonably foreseeable risk of algorithmic*
116 *discrimination.*

117 *"Substantial factor" means a factor that is (i) the principal basis for making a consequential decision, (ii)*
118 *capable of altering the outcome of a consequential decision, and (iii) generated by an artificial intelligence*
119 *system. "Substantial factor" includes any use of an artificial intelligence system to generate any content,*
120 *decision, prediction, or recommendation concerning a consumer that is used as the principal basis to make a*

121 consequential decision concerning the consumer.

122 "Synthetic content" means information, such as images, video, audio clips, and text, that has been
123 significantly modified or generated by algorithms, including by artificial intelligence.

124 "Trade secret" means information, including a formula, pattern, compilation, program, device, method,
125 technique, or process, that (i) derives independent economic value, actual or potential, from not being
126 generally known to, and not being readily ascertainable by proper means by, other persons who can obtain
127 economic value from its disclosure or use and (ii) is the subject of efforts that are reasonable under the
128 circumstances to maintain its secrecy.

129 **§ 59.1-608. Operating standards for developers of high-risk artificial intelligence systems.**

130 A. Each developer of a high-risk artificial intelligence system shall use a reasonable duty of care to
131 protect consumers from any known or reasonably foreseeable risks of algorithmic discrimination. In any
132 enforcement action brought by the Attorney General pursuant to § 59.1-612, there shall be a rebuttable
133 presumption that a developer of a high-risk artificial intelligence system used a reasonable duty of care as
134 required by this subsection if the developer complied with the requirements of this section.

135 B. No developer of a high-risk artificial intelligence system shall offer, sell, lease, give, or otherwise
136 provide to a deployer, integrator, or other developer a high-risk artificial intelligence system unless the
137 developer makes available to the deployer, integrator, or other developer:

138 1. A statement disclosing the intended uses of such high-risk artificial intelligence system;

139 2. Documentation disclosing the following:

140 a. The known or reasonably known limitations of such high-risk artificial intelligence system, including
141 any and all known or reasonably foreseeable risks of algorithmic discrimination arising from the intended
142 uses of such high-risk artificial intelligence system;

143 b. The purpose of such high-risk artificial intelligence system and the intended benefits and uses of such
144 high-risk artificial intelligence system;

145 c. A summary describing how such high-risk artificial intelligence system was evaluated for performance
146 before such high-risk artificial intelligence system was licensed, sold, leased, given, or otherwise made
147 available to a deployer, integrator, or other developer;

148 d. The measures the developer has taken to mitigate reasonable foreseeable risks of algorithmic
149 discrimination that the developer knows arises from deployment or use of such high-risk artificial
150 intelligence system; and

151 e. How an individual can use such high-risk artificial intelligence system and monitor the performance of
152 such high-risk artificial intelligence system for any risk of algorithmic discrimination;

153 3. Documentation including a (i) description of how the high-risk artificial intelligence system was
154 evaluated for performance and for mitigation of algorithmic discrimination before such system was made
155 available to the deployer, integrator, or other developer; (ii) high level summary about how data sources
156 were evaluated for potential bias and appropriate mitigations were applied; (iii) description of the intended
157 outputs of the high-risk artificial intelligence system; (iv) description of the measures the developer has taken
158 to mitigate known or reasonably foreseeable risks of algorithmic discrimination that may arise from the
159 reasonably foreseeable deployment of the high-risk artificial intelligence system; and (v) description of how
160 the high-risk artificial intelligence system should be used, not be used, and be monitored by an individual
161 when such system is used to make, or is a substantial factor in making, a consequential decision; and

162 4. Any additional documentation that is reasonably necessary to assist the deployer, integrator, or other
163 developer in understanding the outputs and monitoring performance of the high-risk artificial intelligence
164 system for risks of algorithmic discrimination.

165 C. Each developer that offers, sells, leases, gives, or otherwise makes available to a deployer, integrator,
166 or other developer a high-risk artificial intelligence system shall make available to the deployer, integrator,
167 or other developer to the extent feasible and necessary, information and documentation through artifacts
168 such as model cards or impact assessments, including any risk management policy designed and implemented
169 and any relevant impact assessment completed, and such documentation and information shall enable the
170 deployer, integrator, other developer, or a third party contracted by the deployer to complete an impact
171 assessment as required in § 59.1-609.

172 D. A developer that also serves as a deployer for any high-risk artificial intelligence system shall not be
173 required to generate the documentation required by this section unless such high-risk artificial intelligence
174 system is provided to an unaffiliated entity acting as a deployer or as otherwise required by law.

175 E. Nothing in this section shall be construed to require a developer to disclose any trade secret, as defined
176 in § 59.1-336, information that could create a security risk, or other confidential or proprietary information.

177 F. High-risk artificial intelligence systems that are in conformity with the latest version of the Artificial
178 Intelligence Risk Management Framework published by the National Institute of Standards and Technology,
179 Standard ISO/IEC 42001 of the International Organization for Standardization, or another nationally or
180 internationally recognized risk management framework for artificial intelligence systems, or parts thereof,
181 shall be presumed to be in conformity with related requirements set out in this section and in associated

182 regulations.

183 G. For any disclosure required pursuant to this section, each developer shall, no later than 90 days after
184 the developer performs an intentional and substantial modification to any high-risk artificial intelligence
185 system, update such disclosure as necessary to ensure that such disclosure remains accurate.

186 H. 1. Each developer of a generative artificial intelligence system that generates or modifies synthetic
187 content shall ensure that the outputs of such high-risk artificial intelligence system (i) are marked and
188 detectable in a manner that is detectable by consumers; (ii) comply with any applicable accessibility
189 requirements, as synthetic content; and (iii) such marking must be applied at the time the output is generated;

190 2. If such synthetic content is an audio, image, or video format that forms part of an evidently artistic,
191 creative, satirical, fictional analogous work or program, such requirement for marking outputs of high-risk
192 artificial intelligence systems pursuant to subdivision 1 shall be limited to a manner that does not hinder the
193 display or enjoyment of such work or program.

194 3. The marking of outputs required by subdivision 1 shall not apply to (i) synthetic content that consists
195 exclusively of text, is published to inform the public on any matter of public interest, or is unlikely to mislead
196 a reasonable person consuming such synthetic content or (ii) the outputs of a high-risk artificial intelligence
197 system that performs an assistive function for standard editing, does not substantially alter the input data
198 provided by the developer, or is used to detect, prevent, investigate, or prosecute any crime as authorized by
199 law.

200 I. Where multiple developers contribute to the development of a high-risk artificial intelligence system,
201 each developer shall be subject to the obligations and operating standards applicable to developers pursuant
202 to this section solely with respect to its activities contributing to the development of the high-risk artificial
203 intelligence system.

204 **§ 59.1-609. Operating standards for deployers of high-risk artificial intelligence systems.**

205 A. Each deployer of a high-risk artificial intelligence system shall use a reasonable duty of care to protect
206 consumers from any known or reasonably foreseeable risks of algorithmic discrimination. In any
207 enforcement action brought by the Attorney General pursuant to § 59.1-612, there shall be a rebuttable
208 presumption that a deployer of a high-risk artificial intelligence system used a reasonable duty of care as
209 required by this subsection if (i) the deployer complied with the provisions of this section and (ii) an
210 independent third party identified by the Attorney General and retained by the developer completed bias and
211 governance audits for the high-risk artificial intelligence system. Not later than July 1, 2026, and at least
212 annually thereafter, the Attorney General shall (a) identify the independent third parties the Attorney General
213 deems qualified to complete bias and governance audits for the purposes of clause (ii) and (b) make a list of
214 such independent third parties available on the Attorney General's website for the high-risk artificial
215 intelligence system.

216 B. No deployer shall deploy or use a high-risk artificial intelligence system to make a consequential
217 decision unless the deployer has designed and implemented a risk management policy and program for such
218 high-risk artificial intelligence system. The risk management policy shall specify the principles, processes,
219 and personnel that the deployer shall use in maintaining the risk management program to identify, mitigate,
220 and document any risk of algorithmic discrimination that is a reasonably foreseeable consequence of
221 deploying or using such high-risk artificial intelligence system to make a consequential decision. Each risk
222 management policy and program designed, implemented, and maintained pursuant to this subsection shall be
223 reasonable considering the guidance and standards set forth in the latest version of:

224 1. The Artificial Intelligence Risk Management Framework published by the National Institute of
225 Standards and Technology;

226 2. Standard ISO/IEC 42001 of the International Organization for Standardization;

227 3. A nationally or internationally recognized risk management framework for artificial intelligence
228 systems with requirements that are substantially equivalent to, and at least as stringent as, the requirements
229 set forth in this section; or

230 4. Any risk management framework for artificial intelligence systems that the Attorney General may
231 designate and is substantially equivalent to, and at least as stringent as, the guidance and standards
232 described in subdivision 1.

233 High-risk artificial intelligence systems that are in conformity with the latest version of the Artificial
234 Intelligence Risk Management Framework published by the National Institute of Standards and Technology,
235 Standard ISO/IEC 42001 of the International Organization for Standardization, or another nationally or
236 internationally recognized risk management framework for artificial intelligence systems, or parts thereof,
237 shall be presumed to be in conformity with related requirements set out in this section and in associated
238 regulations.

239 C. Except as provided in this subsection, no deployer shall deploy or use a high-risk artificial intelligence
240 system to make a consequential decision unless the deployer has completed an impact assessment for such
241 high-risk artificial intelligence system. The deployer shall complete an impact assessment for a high-risk
242 artificial intelligence system (i) before the deployer initially deploys such high-risk artificial intelligence

243 system and (ii) before a significant update to such high-risk artificial intelligence system is made available.

244 Each impact assessment completed pursuant to this subsection shall include, at a minimum:

245 1. A statement by the deployer disclosing (i) the purpose, intended use cases and deployment context of,
246 and benefits afforded by the high-risk artificial intelligence system and (ii) whether the deployment or use of
247 the high-risk artificial intelligence system poses any known or reasonably foreseeable risk of algorithmic
248 discrimination and, if so, (a) the nature of such algorithmic discrimination and (b) the steps that have been
249 taken, to the extent feasible, to mitigate such risk;

250 2. For each post-deployment impact assessment completed pursuant to this subsection, whether the
251 intended use cases of the high-risk artificial intelligence system as updated were consistent with, or varied
252 from, the developer's intended uses of such high-risk artificial intelligence system;

253 3. A description of (i) the categories of data the high-risk artificial intelligence system processes as inputs
254 and (ii) the outputs such high-risk artificial intelligence system produces;

255 4. If the deployer used data to customize the high-risk artificial intelligence system, an overview of the
256 categories of data the deployer used to customize such high-risk artificial intelligence system;

257 5. A list of any metrics used to evaluate the performance and known limitations of the high-risk artificial
258 intelligence system;

259 6. A description of any transparency measures taken concerning the high-risk artificial intelligence
260 system, including any measures taken to disclose to a consumer that such high-risk artificial intelligence
261 system is in use when such high-risk artificial intelligence system is in use;

262 7. A description of any post-deployment monitoring performed and user safeguards provided concerning
263 such high-risk artificial intelligence system, including any oversight process established by the deployer to
264 address issues arising from deployment or use of such high-risk artificial intelligence system as such issues
265 arise; and

266 8. An analysis of such high-risk artificial intelligence system's validity and reliability in accordance with
267 standard industry practices and a description of any metrics used to evaluate the performance and known
268 limitations of such high-risk artificial intelligence system.

269 A single impact assessment may address a comparable set of high-risk artificial intelligence systems
270 deployed or used by a deployer. High-risk artificial intelligence systems that are in conformity with the latest
271 version of the Artificial Intelligence Risk Management Framework published by the National Institute of
272 Standards and Technology, Standard ISO/IEC 42001 of the International Organization for Standardization,
273 or another nationally or internationally recognized risk management framework for artificial intelligence
274 systems, or parts thereof, shall be presumed to be in conformity with related requirements set out in this
275 section and in associated regulations. If a deployer completes an impact assessment for the purpose of
276 complying with another applicable law or regulation, such impact assessment shall be deemed to satisfy the
277 requirements established in this subsection if such impact assessment is reasonably similar in scope and
278 effect to the impact assessment that would otherwise be completed pursuant to this subsection. A deployer
279 that completes an impact assessment pursuant to this subsection shall maintain such impact assessment and
280 all records concerning such impact assessment for three years.

281 Throughout the period of time that a high-risk artificial intelligence system is deployed and for a period of
282 at least three years following the final deployment of such high-risk artificial intelligence system, the
283 deployer shall retain all records concerning each impact assessment conducted on the high-risk artificial
284 intelligence system, including all raw data used to conduct disparate impact testing.

285 D. Not later than the time that a deployer uses a high-risk artificial intelligence system to interact with a
286 consumer, the deployer shall disclose to the consumer that the deployer is interacting with an artificial
287 intelligence system disclosing (i) the purpose of such high-risk artificial intelligence system, (ii) the nature of
288 such system, (iii) the nature of the consequential decision, (iv) the contact information for the deployer, and
289 (v) a description of the artificial intelligence system in plain language of such system, which shall include (a)
290 a description of the personal characteristics or attributes that such system will measure or assess; (b) the
291 method by which the system measures or assesses such attributes or characteristics; (c) how such attributes
292 or characteristics are relevant to the consequential decisions for which the system should be used; (d) any
293 human components of such system; (e) how any automated components of such system are used to inform
294 such consequential decisions; and (f) a direct link to a publicly accessible page on the deployer's website that
295 contains a plain language description of the logic used in the system, including the key parameters that affect
296 the output of the system; the system's outputs; the type and source of data collected from natural persons and
297 processed by the system when it is used to make, or assists in making, a consequential decision; and the
298 results of the most recent impact assessment, or an active link to a webpage where a consumer can review
299 those results.

300 A deployer that has deployed a high-risk artificial intelligence system to make a consequential decision
301 concerning a consumer shall transmit to the consumer the consequential decision within three business days
302 of the consequential decision. If such consequential decision is adverse to such consumer, the deployer shall
303 provide to the consumer (a) a statement disclosing the principal reason or reasons for the consequential

304 decision, including (1) the degree to which and manner in which the high-risk artificial intelligence system
305 contributed to the consequential decision, (2) the type of data that was processed by such system in making
306 the consequential decision, and (3) the sources of such data; (b) pursuant to subdivision A 2 of § 59.1-577,
307 an opportunity to correct any inaccuracies in the consumer's personal data that the high-risk artificial
308 intelligence system processed in making, or as a substantial factor in making, the consequential decision; and
309 (c) an opportunity to appeal such adverse consequential decision concerning the consumer arising from the
310 deployment of such system. Any such appeal shall allow for human review, if technically feasible, unless
311 providing the opportunity for appeal is not in the best interest of the consumer, including instances in which
312 any delay might pose a risk to the life or safety of such consumer.

313 E. Each deployer shall make available, in a manner that is clear and readily available, a statement
314 summarizing how such deployer manages any reasonably foreseeable risk of algorithmic discrimination that
315 may arise from the use or deployment of the high-risk artificial intelligence system.

316 F. For any disclosure required pursuant to this section, each deployer shall, no later than 30 days after
317 the deployer is notified by the developer that the developer has performed an intentional and substantial
318 modification to any high-risk artificial intelligence system, update such disclosure as necessary to ensure that
319 such disclosure remains accurate.

320 G. Any deployer who performs an intentional and substantial modification to any high-risk artificial
321 intelligence system shall comply with the documentation and disclosure requirements for developers pursuant
322 to subsections B through G of § 59.1-608.

323 H. Nothing in this section shall be construed to require a deployer to disclose any trade secret, as defined
324 in § 59.1-336, information that could create a security risk, or other confidential or proprietary information.

325 **§ 59.1-610. Operating standards for integrators of high-risk artificial intelligence systems.**

326 A. On and after September 1, 2026, an integrator shall be considered a developer of a high-risk artificial
327 intelligence system and shall be subject to any obligations and operating standards of a developer pursuant
328 to this chapter with respect to any substantial modifications made by the integrator to such high-risk
329 artificial intelligence system. In such event, the developer that initially placed the high-risk artificial
330 intelligence system on the market shall no longer be considered a developer under this chapter with respect
331 to such substantial modifications made by the integrator.

332 B. On and after September 1, 2026, an integrator shall provide clear and conspicuous notice to the
333 deployer of:

334 1. The name or other identifier of the high-risk artificial intelligence system integrated into a software
335 application provided to the deployer;

336 2. The name and contact information of the developer of the high-risk artificial intelligence system
337 integrated into a software application provided to the deployer;

338 3. If the integrator has adjusted the model weights of the high-risk artificial intelligence system integrated
339 into the software application by exposing it to additional data, a high-level summary of the adjustment
340 process and how that process and the resultant system were evaluated for risk of algorithmic discrimination;

341 4. A high-level summary of any other nonsubstantial modifications made by the integrator; and

342 5. The integrator's applicable policies regarding usage, including any intended limitations or
343 recommended mitigations associated with the intended use of a high-risk artificial intelligence system.

344 C. On and after September 1, 2026, an integrator shall provide to a deployer any information or
345 documentation the developer provides to the integrator pursuant to the developer's obligations under
346 subsections B and C of § 59.1-608, provided that under no circumstance shall the integrator be responsible
347 under this section for any failure to provide documentation to a deployer if that failure was caused by a d
348 eveloper's failure to provide the applicable documentation to the integrator in the first instance.

349 **§ 59.1-611. Exemptions.**

350 A. Nothing in this chapter shall be construed to restrict a developer's or deployer's ability to (i) comply
351 with federal, state, or municipal ordinances or regulations; (ii) comply with a civil, criminal, or regulatory
352 inquiry, investigation, subpoena, or summons by federal, state, local, or other governmental authorities; (iii)
353 cooperate with law-enforcement agencies concerning conduct or activity that the developer or deployer
354 reasonably and in good faith believes may violate federal, state, or local law, ordinances, or regulations; (iv)
355 investigate, establish, exercise, prepare for, or defend legal claims; (v) provide a product or service
356 specifically requested by a consumer; (vi) perform under a contract to which a consumer is a party, including
357 fulfilling the terms of a written warranty; (vii) take steps at the request of a consumer prior to entering into a
358 contract; (viii) take immediate steps to protect an interest that is essential for the life or physical safety of the
359 consumer or another individual; (ix) prevent, detect, protect against, or respond to security incidents, identity
360 theft, fraud, harassment, or malicious or deceptive activities; (x) take actions to prevent, detect, protect
361 against, report, or respond to the production, generation, incorporation, or synthesization of child sex abuse
362 material, or any illegal activity, preserve the integrity or security of systems, or investigate, report, or
363 prosecute those responsible for any such action; (xi) engage in public or peer-reviewed scientific or
364 statistical research in the public interest that adheres to all other applicable ethics and privacy laws and is

365 approved, monitored, and governed by an institutional review board that determines, or similar independent
 366 oversight entities that determine, (a) that the expected benefits of the research outweigh the risks associated
 367 with such research and (b) whether the developer or deployer has implemented reasonable safeguards to
 368 mitigate the risks associated with such research; (xii) assist another developer or deployer with any of the
 369 obligations imposed by this chapter; or (xiii) take any action that is in the public interest in the areas of
 370 public health, community health, or population health, but solely to the extent that such action is subject to
 371 suitable and specific measures to safeguard the public.

372 B. The obligations imposed on developers or deployers by this chapter shall not restrict a developer's or
 373 deployer's ability to (i) conduct internal research to develop, improve, or repair products, services, or
 374 technologies; (ii) effectuate a product recall; (iii) identify and repair technical errors that impair existing or
 375 intended functionality; or (iv) perform internal operations that are reasonably aligned with the expectations
 376 of the consumer or reasonably anticipated based on the consumer's existing relationship with the developer
 377 or deployer.

378 C. Nothing in this chapter shall be construed to impose any obligation on a developer or deployer to
 379 disclose trade secrets or information protected from disclosure by state or federal law.

380 D. The obligations imposed on developers or deployers by this chapter shall not apply where compliance
 381 by the developer or deployer with such obligations would violate an evidentiary privilege under federal law
 382 or the laws of the Commonwealth.

383 E. Nothing in this chapter shall be construed to impose any obligation on a developer or deployer that
 384 adversely affects the legally protected rights or freedoms of any person, including the rights of any person to
 385 freedom of speech or freedom of the press guaranteed in the First Amendment to the Constitution of the
 386 United States or under the Virginia Human Rights Act (§ 2.2-3900 et seq.).

387 F. The obligations imposed on developers or deployers by this chapter shall not apply to any artificial
 388 intelligence system that is acquired by or for the federal government or any federal agency or department,
 389 including the U.S. Department of Commerce, the U.S. Department of Defense, and the National Aeronautics
 390 and Space Administration, unless such artificial intelligence system is a high-risk artificial intelligence
 391 system that is used to make, or is a substantial factor in making, a decision concerning employment or
 392 housing.

393 G. For the purposes of this subsection:

394 "Affiliate" means the same as that term is defined in § 6.2-899.

395 "Bank" means the same as that term is defined in § 6.2-800.

396 "Credit union" means the same as that term is defined in § 6.2-1300.

397 "Federal credit union" means a credit union duly organized under federal law.

398 "Mortgage lender" means the same as that term is defined in § 6.2-1600.

399 "Out-of-state bank" means the same as that term is defined in § 6.2-836.

400 "Out-of-state credit union" means a credit union organized and doing business in another state.

401 "Savings institution" means the same as that term is defined in § 6.2-1100.

402 "Subsidiary" means the same as that term is defined in § 6.2-700.

403 The obligations imposed on developers or deployers by this chapter shall be deemed satisfied for any
 404 bank, out-of-state bank, credit union, federal credit union, mortgage lender, out-of-state credit union, savings
 405 institution, or any affiliate or subsidiary thereof if such bank, out-of-state bank, credit union, federal credit
 406 union, mortgage lender, out-of-state credit union, savings institution, or affiliate or subsidiary is subject to
 407 the jurisdiction of any state or federal regulator under any published guidance or regulations that apply to
 408 the use of high-risk artificial intelligence systems and such guidance or regulations.

409 H. For purposes of this subsection, "insurer" means the same as that term is defined in § 38.2-100.

410 The provisions of this chapter shall not apply to any insurer, or any high-risk artificial intelligence system
 411 developed or deployed by an insurer for use in the business of insurance, if such insurer is regulated and
 412 supervised by the State Corporation Commission or a comparable federal regulating body and subject to
 413 examination by such entity under any existing statutes, rules, or regulations pertaining to unfair trade
 414 practices and unfair discrimination prohibited under Chapter 5 (§ 38.2-500 et seq.) of Title 38.2, or
 415 published guidance or regulations that apply to the use of high-risk artificial intelligence systems and such
 416 guidance or regulations aid in the prevention and mitigation of algorithmic discrimination caused by the use
 417 of a high-risk artificial intelligence system or any risk of algorithmic discrimination that is reasonably
 418 foreseeable as a result of the use of a high-risk artificial intelligence system. Nothing in this chapter shall be
 419 construed to delegate existing regulatory oversight of the business of insurance to any department or agency
 420 other than the Bureau of Insurance of the Virginia State Corporation Commission.

421 I. The provisions of this chapter shall not apply to the development of an artificial intelligence system that
 422 is used exclusively for research, training, testing, or other pre-deployment activities performed by active
 423 participants of any sandbox software or sandbox environment established and subject to oversight by a
 424 designated agency or other government entity and that is in compliance with the provisions of this chapter.

425 J. The provisions of this chapter shall not apply to a developer or deployer, or other person who develops,

426 *deploys, puts into service, or intentionally modifies, as applicable, a high-risk artificial intelligence system*
427 *that (i) has been approved, authorized, certified, cleared, developed, or granted by a federal agency acting*
428 *within the scope of the federal agency's authority, or by a regulated entity subject to the supervision and*
429 *regulation of the Federal Housing Finance Agency or (ii) is in compliance with standards established by a*
430 *federal agency or by a regulated entity subject to the supervision and regulation of the Federal Housing F*
431 *inance Agency, if the standards are substantially equivalent or more stringent than the requirements of this*
432 *chapter.*

433 *K. The provisions of this chapter shall not apply to a developer or deployer, or other person that (i)*
434 *facilitates or engages in the provision of telehealth services, as defined in § 32.1-122.03:1, or (ii) is a*
435 *covered entity within the meaning of the federal Health Insurance Portability and Accountability Act of 1996*
436 *(42 U.S.C. § 1320d et seq.) and the regulations promulgated under such federal act, as both may be amended*
437 *from time to time, and is providing (a) health care recommendations that (1) are generated by an artificial*
438 *intelligence system and (2) require a health care provider, as defined in § 8.01-581.1, to take action to*
439 *implement the recommendations or (b) services utilizing an artificial intelligence system for an*
440 *administrative, quality measurement, security, or internal cost or performance improvement function.*

441 *L. If a developer or deployer engages in any action authorized by an exemption set forth in this section,*
442 *the developer or deployer bears the burden of demonstrating that such action qualifies for such exemption.*

443 *M. If a developer or deployer withholds information pursuant to an exemption set forth in this chapter for*
444 *which disclosure would otherwise be required by this chapter, including the exemption from disclosure of*
445 *trade secrets, the developer or deployer shall notify the subject of disclosure and provide a basis for*
446 *withholding the information. If a developer or deployer redacts any information pursuant to an exemption*
447 *from disclosure, the developer or deployer shall notify the subject of disclosure that the developer or*
448 *deployer is redacting such information and provide the basis for such decision to redact.*

449 **§ 59.1-612. Enforcement; civil penalties.**

450 *A. The Attorney General shall have exclusive authority to enforce the provisions of this chapter.*

451 *B. Whenever the Attorney General has reasonable cause to believe that any person has engaged in or is*
452 *engaging in any violation of this chapter, the Attorney General is empowered to issue a civil investigative*
453 *demand. The provisions of § 59.1-9.10 shall apply mutatis mutandis to civil investigative demands issued*
454 *pursuant to this section. In rendering and furnishing any information requested pursuant to a civil*
455 *investigative demand issued pursuant to this section, a developer, integrator, or deployer may redact or omit*
456 *any trade secrets or information protected from disclosure by state or federal law. If a developer, integrator,*
457 *or deployer refuses to disclose, redacts, or omits information based on the exemption from disclosure of trade*
458 *secrets, such developer, integrator, or deployer shall affirmatively state to the Attorney General that the basis*
459 *for nondisclosure, redaction, or omission is because such information is a trade secret. To the extent that any*
460 *information requested pursuant to a civil investigative demand issued pursuant to this section is subject to*
461 *attorney-client privilege or work-product protection, disclosure of such information pursuant to the civil*
462 *investigative demand shall not constitute a waiver of such privilege or protection. Any information, statement*
463 *, or documentation provided to the Attorney General pursuant to this section shall be exempt from disclosure*
464 *under the Virginia Freedom of Information Act (§ 2.2-3700 et seq.).*

465 *C. Notwithstanding any contrary provision of law, the Attorney General may cause an action to be*
466 *brought in the appropriate circuit court in the name of the Commonwealth to enjoin any violation of this*
467 *chapter. The circuit court having jurisdiction may enjoin such violation notwithstanding the existence of an*
468 *adequate alternative remedy at law.*

469 *D. Any person who violates the provisions of this chapter shall be subject to a civil penalty in an amount*
470 *not to exceed \$1,000 plus reasonable attorney fees, expenses, and costs, as determined by the court. Any*
471 *person who willfully violates the provisions of this chapter shall be subject to a civil penalty in an amount not*
472 *less than \$1,000 and not more than \$10,000 plus reasonable attorney fees, expenses, and costs, as*
473 *determined by the court. Such civil penalties shall be paid into the Literary Fund.*

474 *E. Each violation of this chapter shall constitute a separate violation and shall be subject to any civil*
475 *penalties imposed under this section.*

476 *F. The Attorney General may require that a developer disclose to the Attorney General any statement or*
477 *documentation described in this chapter if such statement or documentation is relevant to an investigation*
478 *conducted by the Attorney General. The Attorney General may also require that a deployer disclose to the*
479 *Attorney General any risk management policy designed and implemented, impact assessment completed, or*
480 *record maintained pursuant to this chapter if such risk management policy, impact assessment, or record is*
481 *relevant to an investigation conducted by the Attorney General.*

482 *G. In an action brought by the Attorney General pursuant to this section, it shall be an affirmative defense*
483 *that the developer, integrator, or deployer (i) discovers a violation of any provision of this chapter through*
484 *red-teaming; (ii) no later than 45 days after discovering such violation (a) cures such violation and (b)*
485 *provides notice to the Attorney General in a form and manner as prescribed by the Attorney General that*
486 *such violation has been cured and evidence that any harm caused by such violation has been mitigated; and*

487 (iii) is otherwise in compliance with the requirements of this chapter.

488 H. Prior to causing an action against a developer, integrator, or deployer for a violation of this chapter
489 pursuant to subsection C, the Attorney General shall determine, in consultation with the developer,
490 integrator, or deployer, if it is possible to cure the violation. If it is possible to cure such violation, the
491 Attorney General may issue a notice of violation to the developer, integrator, or deployer and afford the
492 developer, integrator, or deployer the opportunity to cure such violation within 45 days of the receipt of such
493 notice of violation. In determining whether to grant such opportunity to cure such violation, the Attorney G
494 eneral shall consider (i) the number of violations; (ii) the size and complexity of the developer, integrator, or
495 deployer; (iii) the nature and extent of the developer's, integrator's, or deployer's business; (iv) the
496 substantial likelihood of injury to the public; (v) the safety of persons or property; and (vi) whether such
497 violation was likely caused by human or technical error. If the developer, integrator, or deployer fails to cure
498 such violation within 45 days of the receipt of such notice of violation, the Attorney General may proceed
499 with such action.

500 I. Nothing in this chapter shall create a private cause of action in favor of any person aggrieved by a
501 violation of this chapter.

502 **§ 59.1-613. Construction of chapter.**

503 A. This chapter is declared to be remedial, with the purposes of protecting consumers and ensuring
504 consumers receive information about consequential decisions affecting them. The provisions of this chapter
505 granting rights or protections to consumers shall be construed broadly and exemptions construed narrowly.

506 B. If any provision of this chapter or its application to any person or circumstance is held invalid, the
507 invalidity shall not affect other provisions or applications of this chapter that can be given effect without the
508 invalid provision or application, and to this end all the provisions of this chapter are hereby expressly
509 declared to be severable.

510 **2. That the provisions of this act shall become effective on July 1, 2026.**

511 **3. That compliance with the provisions of Chapter 58 (§ 59.1-607 et seq.) of Title 59.1 of the Code of**
512 **Virginia, as created by this act, shall not (i) relieve a person from liability for any causes of action that**
513 **existed at common law or by statute prior to July 1, 2026, or (ii) be construed to modify or otherwise**
514 **affect, preempt, limit, or displace any causes of action that existed at common law or by statute prior to**
515 **July 1, 2026.**