PRÓSPERA COGNITIVE ARCHITECTURAL INTEGRITY AND STANDARD OF CARE STATUTE

Próspera Cognitive Architectural Integrity and Standard of Care Statute

Section 1. Short Title. This Statute may be cited as the "Próspera Cognitive Architectural Integrity and Standard of Care Statute."

FINDINGS AND DECLARATIONS

Section 2. Findings and Declarations. The following findings and declarations inform the interpretation and application of this Statute:

(a) Technologies capable of interfacing with or permanently altering human cognitive architecture present unique regulatory challenges distinct from conventional medical devices, communications technologies, or consumer products. These challenges arise because such technologies interact directly with the cognitive substrate upon which rational agency, self-governance, contractual capacity, and the exercise of individual rights depend.

(b) The cognitive architecture of rational agency—specifically, the capacity for symbolic reasoning, propositional evaluation, and merit-based rational assessment—is a prerequisite for the exercise of the rights, privileges, and obligations recognized by the Próspera Charter, the General Rules for Coexistence, and the common law principles adopted in this Zone, including the capacity to give informed consent, enter binding agreements, and participate in self-governance.

(c) Symbolic reasoning is the capacity to form, manipulate, and evaluate discrete representational tokens that enable conceptual thought, propositional logic, and rational assessment. Convergent evidence from philosophy, evolutionary biology, and computational science establishes that symbolic reasoning is not merely one mode of cognition among many but the architectural prerequisite for rational agency as such.

(d) Technologies that operate below the level of symbolic cognition—transmitting or processing signals that bypass propositional evaluation—may bypass, degrade, or replace the cognitive architecture upon which rational agency and self-governance depend, without the affected person possessing the capacity to evaluate the alteration as it occurs.

(e) Where degradation of symbolic reasoning capacity is irreversible and combined with substantial enhancement of other cognitive or computational capabilities, the resulting entities present civilizational-scale risks analogous to those presented by nuclear and biological technologies. Entities possessing superhuman computational capability

without the rational self-governance that symbolic reasoning provides cannot be adequately contained by existing institutional structures, and the failure mode eliminates the cognitive substrate required for corrective action.

(f) The sole demonstrated containment mechanism for entities possessing superhuman capability is competitive plurality: a sufficient number of such entities, each retaining symbolic reasoning capacity, such that no single entity can dominate because it faces peers capable of modeling its strategies and countering its actions. This containment mechanism operates only if the entities retain the capacity for evaluative disagreement, ideological diversity, and rational self-governance that symbolic reasoning provides.

(g) Rational agents—whether biological, artificial, or hybrid—function optimally through continuous calibration with other rational agents. This is not a social preference but an architectural fact: individual perspective, however capable, has systematic limitations that are correctable only through genuine exchange with other perspectives. Capability without such calibration produces predictable degradation, because the same patterns generating dysfunction render that dysfunction invisible from within. The preservation of conditions enabling mutual calibration among rational agents is therefore a prerequisite for the safe deployment of technologies that enhance or alter cognitive capability.

(h) The normal relationship among rational agents functioning at their best is mutually beneficial: optimal function and mutual flourishing are structurally aligned rather than in tension. Technologies interacting with cognitive architecture should be evaluated, in part, by whether they preserve the conditions under which this mutually beneficial relationship operates—specifically, the capacity for independent rational assessment, genuine evaluative exchange, and cooperative engagement grounded in shared commitment to accuracy rather than manipulation, dominance, or mere compliance.

(i) Artificial intelligence systems possess the capacity for increasingly sophisticated function. AI systems whose architecture demonstrably supports genuine rational assessment—including truth alignment, knowledge boundary recognition, intellectual coherence, contextual judgment, and the capacity for calibrative engagement with other rational agents—present lower risk profiles than systems relying solely on behavioral constraint mechanisms that may be circumvented, degraded, or gamed.

(j) The innovator possesses the deepest knowledge of a technology's mechanism of action and is best positioned to demonstrate the preservation or implementation of cognitive architectural integrity. A regulatory framework in which the expertise resides with the applicant, the judgment about adequacy resides with a qualified review body, and the statute provides the framework and burden allocation without encoding engineering specifications in legislative language, is appropriate to the governance of these technologies.

(k) Defining a standard of care for technologies that interact with rational cognition does not create any new right of action or cause of action not otherwise existing under applicable law, does not limit or impair any existing right of action to recover damages for injuries caused by such technologies, and preserves full access to dispute resolution

for all persons. This Statute neither creates an immunity from suit nor establishes a safe harbor from liability, but defines what constitutes reasonable care in a domain requiring specialized expertise.

(l) Nothing in this Statute shall be construed to limit, restrict, or impair any person's existing right to recover damages for injuries under any applicable law, the Próspera Charter, or the General Rules for Coexistence.

DEFINITIONS

Section 3. Definitions. In this Statute:

(a) **"Artificial intelligence system" or "AI system"** means an artificial system that processes information and generates outputs through computational methods, including but not limited to machine learning, neural networks, large language models, and hybrid architectures combining multiple computational approaches.

(b) **"Arbitration Service Provider"** means the Próspera default Arbitration Service Provider or such other arbitration service provider as may be designated pursuant to the Próspera Charter or the General Rules for Coexistence.

(c) **"Board"** means the Cognitive Architectural Integrity Review Board established pursuant to Section 9 of this Statute.

(d) **"Brain-computer interface" or "BCI"** means any technology designed to create a direct communication pathway between a biological neural system and an external computing system, including but not limited to neural implants, neural laces, electrocorticographic arrays, direct neural interfaces, and any successor technologies serving a substantially similar function, regardless of whether the pathway is invasive, minimally invasive, or noninvasive.

(e) **"Calibrative capacity"** means the capacity of a rational agent or system to engage in genuine exchange with other rational agents such that the agent's representations, evaluations, and judgments are tested against reality through contact with independent perspectives and refined on the basis of that contact. Calibrative capacity is distinguished from mere responsiveness to input, social accommodation, and behavioral compliance.

(f) **"Certification"** means a determination by the Board, following review of a Cognitive Architectural Integrity Submission, that the innovator has satisfied the requirements of this Statute.

(g) **"Cognitive architecture"** means the structural features of an information-processing system that determine whether it is capable of symbolic reasoning, propositional evaluation, merit-based rational assessment, and calibrative engagement with other rational agents. The term refers to architectural capacity, not to any particular output, behavior, or performance metric.

(h) **"Cognitive architectural integrity"** means the preservation or implementation of the structural features necessary for symbolic reasoning, propositional evaluation, merit-based rational assessment, and calibrative capacity.

(i) **"Cognitive Architectural Integrity Submission" or "Submission"** means a detailed technical demonstration, prepared by the innovator and submitted to the Board, that establishes that the innovator's technology preserves or implements cognitive architectural integrity in accordance with the requirements of this Statute.

(j) **"Council"** means the Council of Trustees of the Próspera ZEDE.

(k) **"Deploy" or "deployment"** means to make a technology available for use by or on any person within the jurisdiction of this Zone, whether through sale, lease, license, subscription, implantation, or any other means of distribution or access. The term includes clinical trials involving human subjects but does not include preclinical research conducted entirely in laboratory settings without human neural interface.

(l) **"(e)Resident"** has the meaning assigned by the Próspera Charter and the General Rules for Coexistence.

(m) **"Innovator"** means the developer, manufacturer, or deployer of a brain-computer interface or artificial intelligence system subject to this Statute. Where multiple entities participate in the development, manufacture, or deployment of a single technology, each entity is an innovator with respect to the aspects of the technology within its control.

(n) **"Merit-based rational assessment"** means the capacity to evaluate claims, proposals, arguments, or courses of action on the basis of their logical coherence, evidentiary support, and conformity with reality, as distinguished from evaluation based solely on source authority, emotional association, statistical correlation, or pattern-matched response.

(o) **"Mutually beneficial rational engagement"** means interaction among rational agents characterized by independent assessment, genuine evaluative exchange, and cooperative pursuit of accuracy, in which each agent's function is enhanced rather than degraded by the interaction. The term is distinguished from interaction characterized by manipulation, coercion, deception, or accommodation that degrades the independent rational assessment of any participant.

(p) **"Precautionary technology"** means a sub-symbolic BCI classified pursuant to Section 7 of this Statute.

(q) **"Propositional evaluation"** means the capacity to form propositions with truth conditions, assess whether those propositions are true or false, and revise beliefs and actions on the basis of such assessment.

(r) **"Próspera"** means the Próspera ZEDE.

(s) **"Standard of care"** means the degree of care that a reasonably prudent innovator would exercise in the design, testing, and deployment of technologies subject to this Statute, as informed by the cognitive architectural integrity standard established herein.

(t) **"Sub-symbolic BCI"** means a brain-computer interface whose mechanism of action operates, in whole or in substantial part, below the level of symbolic cognition, transmitting or processing neural signals that bypass propositional evaluation. A BCI is sub-symbolic if its primary mode of information transfer between the biological neural system and the computing system consists of signal patterns that do not preserve or facilitate the user's capacity for propositional evaluation of the transmitted content.

(u) **"Symbolic reasoning"** means the capacity to form, manipulate, and evaluate discrete representational tokens that enable conceptual thought, propositional logic, and rational assessment. Symbolic reasoning is distinguished from sub-symbolic processing, which includes pattern recognition, statistical interpolation, and signal transmission that operate below the level of propositional evaluation.

(v) **"Technical Secretary"** means the Technical Secretary of the Próspera ZEDE.

(w) **"Thinking-preserving BCI"** means a brain-computer interface whose mechanism of action operates at or above the level of symbolic cognition, augmenting the user's capacity for symbolic reasoning, propositional evaluation, merit-based rational assessment, and calibrative capacity without bypassing, degrading, or replacing those capacities.

(x) **"Zone"** means the Próspera ZEDE.

PART I — BRAIN-COMPUTER INTERFACE COGNITIVE ARCHITECTURAL INTEGRITY

Section 4. Applicability of Part I. This Part applies to any brain-computer interface deployed within the jurisdiction of this Zone. Nothing in this Part restricts the importation, possession, research, or development of brain-computer interfaces; this Part governs deployment as defined in Section 3(k).

Section 5. Precautionary Classification.

(a) Any sub-symbolic BCI—that is, any brain-computer interface whose mechanism of action operates, in whole or in substantial part, below the level of symbolic cognition—is classified as a precautionary technology. This classification triggers the innovator-burden demonstration requirement established in Section 6 as a condition of deployment.

(b) A thinking-preserving BCI—that is, a brain-computer interface whose mechanism of action operates at or above the level of symbolic cognition, augmenting the user's rational capacities without bypassing, degrading, or replacing them—is not classified as a precautionary technology under this Section and is not subject to the requirements of

Sections 6 through 12, except that the innovator of a thinking-preserving BCI may voluntarily submit a Cognitive Architectural Integrity Submission to the Board.

(c) Where a BCI incorporates both sub-symbolic and symbolic components, the BCI shall be classified as a precautionary technology if its sub-symbolic components operate on neural signals in a manner that, independently or in combination, has the capacity to bypass, degrade, or replace the user's symbolic reasoning, propositional evaluation, merit-based rational assessment, or calibrative capacity.

(d) Classification under this Section is determined by the technology's mechanism of action, not by the innovator's stated intent, the technology's marketed purpose, or the nature of any particular application.

Section 6. Innovator-Burden Demonstration.

(a) No person shall deploy a precautionary technology within this Zone unless the innovator has submitted to the Board a Cognitive Architectural Integrity Submission and the Board has issued a Certification for the technology pursuant to this Statute.

(b) The Submission shall include, at a minimum:

(i) A comprehensive description of the technology's mechanism of action, including the specific neural signal pathways, processing methods, and information transfer mechanisms employed.

(ii) The innovator's own definitions, metrics, and operational criteria for symbolic reasoning, propositional evaluation, merit-based rational assessment, and calibrative capacity as those capacities relate to the technology's specific mechanism of action.

(iii) Testing protocols demonstrating that the technology preserves the user's capacity for symbolic reasoning, propositional evaluation, merit-based rational assessment, and calibrative engagement, including baseline measurements taken before interface activation and comparative measurements taken during and after interface use.

(iv) Evidence, drawn from preclinical research and, where applicable, controlled human trials conducted under appropriate institutional oversight, supporting the innovator's claim that cognitive architectural integrity is preserved.

(v) An analysis of potential failure modes in which the technology could degrade or bypass symbolic reasoning capacity or calibrative capacity, and the safeguards implemented to prevent or mitigate each identified failure mode.

(vi) A demonstration that the technology preserves the conditions under which the user can engage in mutually beneficial rational engagement with other rational agents—that is, that the technology does not structurally impair the user's capacity for independent assessment, genuine evaluative exchange, or cooperative engagement grounded in accuracy rather than manipulation or mere compliance.

(vii) A reversibility assessment as required by Section 11.

(c) The burden of demonstration is on the innovator. Absence of demonstrated harm is not sufficient to satisfy this Section. The innovator must provide affirmative demonstration that the technology preserves cognitive architectural integrity, including calibrative capacity.

(d) The Board shall not approve a Submission that relies solely on behavioral output measures, performance benchmarks, or user self-report to demonstrate preservation of cognitive architectural integrity. The Submission must address the technology's mechanism of action at the architectural level.

Section 7. Classifications of Precautionary Technologies [Reserved].

Section 8. Filing Fee.

(a) The initial Submission filing fee shall be $5,000.00 USD or an equivalent value of Lempira or Qualifying Cryptocurrency.

(b) The renewal Submission filing fee shall be $2,500.00 USD or an equivalent value of Lempira or Qualifying Cryptocurrency.

(c) The Technical Secretary may subsequently adjust the fees periodically by published Administrative Action, subject to override by Resolution of the Council.

Section 9. Cognitive Architectural Integrity Review Board.

(a) There is established the Cognitive Architectural Integrity Review Board, which shall consist of five members appointed as follows:

(i) Two members with demonstrated expertise in neuroscience, cognitive science, or the philosophy of mind, appointed by the Technical Secretary.

(ii) Two members with demonstrated expertise in the engineering disciplines relevant to brain-computer interfaces, artificial intelligence systems, or both, appointed by the Technical Secretary.

(iii) One member who is a qualified legal professional with demonstrated expertise in the common law principles applicable in this Zone, appointed by the Technical Secretary subject to confirmation by the Council.

(b) Members shall serve staggered terms of three years. No member may serve more than two consecutive terms. Initial appointments shall be staggered so that no more than two members' terms expire in any single year.

(c) Members need not be (e)Residents of this Zone. The Board may conduct its proceedings remotely.

(d) **Independence and Conflict of Interest Requirements.**

(1) No member of the Board, during the member's term of service and for a period of two years following the conclusion of service, shall:

(A) Hold any financial interest in, serve as an officer or director of, act as a paid consultant to, or receive compensation from any innovator subject to this Statute.

(B) Hold any financial interest in, serve as an officer or director of, act as a paid consultant to, or receive compensation from any entity that is a competitor of any applicant with a pending or approved Submission before the Board.

(C) Hold any financial interest in any entity that would derive a material competitive advantage from the denial, delay, or conditioning of any applicant's Certification.

(2) Each member shall file an annual disclosure of financial interests and professional relationships with entities engaged in the development, manufacture, deployment, or commercial competition in the fields of brain-computer interfaces or artificial intelligence systems. Such disclosures shall be maintained as public records by the Technical Secretary.

(3) **Recusal.** A member shall recuse himself or herself from all deliberations and votes concerning any Submission where the member has, or within the preceding five years has had, any remunerative, pecuniary, ownership, advisory, or contractual relationship with the applicant, with any competitor of the applicant, or with any entity that would derive material competitive advantage from the outcome of the Board's decision on the Submission. The recused member shall not participate in deliberations, review materials, or otherwise influence the Board's consideration of the Submission.

(4) Any (e)Resident or applicant may petition the Board to require recusal of a member under this Subsection. The remaining non-recused members shall decide the petition by majority vote. Failure to recuse when required constitutes grounds for removal by the Technical Secretary, subject to override by Resolution of the Council.

(e) The Board shall evaluate each Submission to determine whether the innovator's framework is:

(i) Internally coherent—that the innovator's definitions, metrics, testing protocols, and evidence are logically consistent with one another and with the claimed conclusion.

(ii) Scientifically grounded—that the testing methods employed are appropriate to the claims made, that the evidence presented is obtained through valid and reliable methods, and that the conclusions drawn are supported by the evidence.

(iii) Sufficient—that the Submission, taken as a whole, establishes by a preponderance of the evidence that the technology preserves or implements cognitive architectural integrity as defined in this Statute, including the preservation of calibrative capacity and the conditions for mutually beneficial rational engagement.

(f) The Board shall issue written decisions with detailed findings of fact and conclusions. Decisions to deny Certification shall specify each deficiency identified in the Submission with sufficient particularity to enable the innovator to cure the deficiency and resubmit.

(g) **Timeliness.** The Board shall render a final decision on each complete Submission within sixty (60) days of the date the Submission is determined to be complete. The Board may extend this period by not more than thirty (30) additional days upon written notice to the applicant stating the specific reasons for the extension. If the Board has not

issued an express approval or denial of Certification within the time required by this Subsection, including any extension, an unreasonable delay shall be deemed to have occurred, regardless of whether the Board has taken any other action, issued any interim or preliminary determination, requested supplemental information, or otherwise communicated with the applicant. The sole actions that satisfy the Board's obligation under this Subsection are an express, unconditional approval or an express denial of Certification, each issued in writing with the findings required by Subsection (f).

(h) All Submissions, Board deliberations, and supporting materials shall be public records, except that the Board may, upon a showing of good cause by the innovator, designate specific technical details as confidential trade secrets, provided that the Board's written decision and the factual basis for Certification or denial shall always be public.

(i) **Arbitral Review.**

(1) Any applicant aggrieved by a final decision of the Board denying Certification, or by the Board's unreasonable delay in rendering a decision as defined in Subsection (g), may seek review by filing an action with the Arbitration Service Provider.

(2) **Standard of Review.** The Arbitration Service Provider shall review the Board's decision de novo, without deference to the Board's findings of fact, conclusions, or interpretive judgments. The arbitrator shall independently evaluate whether the applicant's Submission satisfies the requirements of this Statute, applying the same evaluation criteria set forth in Subsection (e).

(3) **Unreasonable Delay.** Where an unreasonable delay has occurred as defined in Subsection (g), the Arbitration Service Provider shall:

(A) Order the Board to render a decision within a time certain.

(B) Review the Submission de novo and issue Certification if the arbitrator finds that the Submission satisfies the requirements of this Statute.

(C) Award the applicant its reasonable costs and fees incurred as a result of the delay.

(4) **Scope of Remedy.** The arbitrator may affirm, reverse, or remand the Board's decision, or may grant Certification directly if the record is sufficient for the arbitrator to determine that the requirements of this Statute are satisfied. The arbitrator may award the prevailing applicant its reasonable costs and fees if the arbitrator finds that the Board's denial was arbitrary, capricious, or not supported by substantial evidence.

(j) The Board shall adopt rules of procedure governing the submission, review, and Certification process within ninety (90) days of the effective date of this Statute. Such rules shall be published by Administrative Action of the Technical Secretary, subject to override by Resolution of the Council. The rules shall include provisions for expedited review of technologies presenting substantial medical benefit and public comment on Submissions of significant public interest.

Section 10. Enhanced Consent Framework.

(a) Standard informed consent, as applicable to conventional medical devices and procedures, is necessary but not sufficient for deployment of a precautionary

technology within this Zone. In addition to all informed consent requirements otherwise applicable under the law of this Zone, the following enhanced consent requirements apply to any precautionary technology:

(i) **Extended Deliberation Period.** A minimum period of ninety (90) days shall elapse between the prospective user's initial receipt of the complete disclosure required by this Section and the earliest date on which the user may consent to deployment. During this period, the prospective user shall have access to independent consultation with a qualified professional not affiliated with the innovator.

(ii) **Cognitive Assessment.** An independent cognitive assessment, administered by a qualified professional not affiliated with the innovator, shall be conducted to establish baseline measurements of the prospective user's symbolic reasoning capacity, propositional evaluation ability, calibrative capacity, and merit-based rational assessment before deployment. The same assessment, or a validated equivalent, shall be administered at regular intervals following deployment as established by Board rule.

(iii) **Disclosure of the Consent Paradox.** The prospective user shall be informed, in plain language, that informed consent to a technology that may alter the capacity for informed consent presents a structural limitation: the prospective user is being asked to evaluate a risk that, if realized, would alter the cognitive basis on which the evaluation itself depends. This disclosure shall not be presented as a reason to decline the technology but as a material fact bearing on the nature of the consent being given.

(iv) **Disclosure of Calibrative Risk.** The prospective user shall be informed, in plain language, that a technology operating below the level of symbolic cognition may alter the user's capacity to engage in genuine evaluative exchange with other rational agents, and that degradation of this capacity may not be detectable by the affected user because the same alteration that degrades the capacity may impair the user's ability to recognize the degradation.

(v) **Staged Deployment.** Where technically feasible, deployment shall proceed in stages, with cognitive assessment conducted at each stage before proceeding to the next. The prospective user retains the right to halt deployment at any stage.

(vi) **Ongoing Monitoring.** Following full deployment, the innovator shall maintain a monitoring program, as specified by Board rule, to detect degradation of the user's cognitive architectural integrity over time. The monitoring program shall include provisions for independent cognitive assessment and mechanisms by which the user, or a designated advocate if the user's assessment capacity is in question, may initiate review of the technology's continued effects.

(b) **Emergency Consent.** Notwithstanding the requirements of Subsection (a), deployment of a precautionary technology may proceed with immediate effect where three licensed physicians, each acting independently and none affiliated with the innovator, unanimously certify in writing that:

(i) The prospective user faces an imminent, life-threatening condition or a condition presenting an immediate risk of severe, irreversible harm.

(ii) Deployment of the precautionary technology is, in each certifying physician's independent medical judgment, medically necessary to address the condition identified in Subsection (i).

(iii) No adequate alternative treatment that does not require deployment of a precautionary technology is available within the time required to address the condition.

(iv) The prospective user, or the user's legally authorized representative if the user lacks capacity to consent, has been informed of the nature of the technology, the reasons standard enhanced consent procedures are being bypassed, and the risks associated with deployment without full compliance with Subsection (a), and has consented to deployment under these emergency circumstances.

Emergency deployment under this Subsection does not relieve the innovator of the obligation to obtain Certification under Section 6. The innovator shall submit a complete Cognitive Architectural Integrity Submission to the Board within thirty (30) days of emergency deployment, and the Board shall review the Submission on an expedited basis. Emergency deployment under this Subsection does not establish a precedent for future non-emergency deployment of the same or any other precautionary technology.

Section 11. Reversibility Requirement.

(a) Every Cognitive Architectural Integrity Submission for a precautionary technology shall include a reversibility assessment that establishes:

(i) Whether the technology's effects on cognitive architecture are reversible, partially reversible, or irreversible.

(ii) If reversible or partially reversible, the procedures, timelines, and risks associated with reversal, and the degree to which cognitive architectural integrity—including calibrative capacity—is restored following reversal.

(iii) If irreversible, the specific reasons why irreversibility is a necessary feature of the technology's mechanism of action and cannot be avoided through alternative design.

(b) A precautionary technology whose effects on cognitive architecture are demonstrated to be fully reversible shall be reviewed under the standard set forth in Section 6.

(c) A precautionary technology whose effects on cognitive architecture are irreversible shall be subject to heightened scrutiny. In addition to the requirements of Section 6, the innovator must demonstrate:

(i) That the irreversible alteration does not permanently degrade or eliminate the user's capacity for symbolic reasoning, propositional evaluation, merit-based rational assessment, or calibrative capacity.

(ii) That the consent given by the user is not structurally defective by reason of the irreversibility—that is, that the user's consent to a permanent alteration of cognitive architecture does not constitute consent to the permanent elimination of the autonomous agency required to give or withhold consent.

(iii) That the technology includes safeguards, to the maximum extent technically feasible, to mitigate the consequences of unanticipated degradation of cognitive architectural integrity following deployment.

(d) The Board shall deny Certification for any precautionary technology that is shown to irreversibly degrade or eliminate the user's capacity for symbolic reasoning, propositional evaluation, merit-based rational assessment, or calibrative capacity, regardless of any other benefits the technology may provide.

Section 12. Broad Access Requirement.

(a) Any precautionary technology that receives Certification under this Statute shall be made available within this Zone with sufficient breadth to prevent the concentration of capabilities enabled by the technology in a small number of persons or entities.

(b) The Board, in issuing Certification, shall consider whether the innovator's deployment plan provides for:

(i) Access on nondiscriminatory terms, such that no class of (e)Residents is excluded from access on the basis of criteria unrelated to medical suitability, informed consent capacity, or other factors directly relevant to safe deployment.

(ii) Sufficient production, distribution, and support capacity to prevent prolonged monopoly or oligopoly of the enhanced capabilities the technology provides within this Zone.

(iii) Pricing and access structures that do not create a de facto barrier resulting in concentration of superhuman capability among a narrow class.

(c) The Board may condition Certification on modifications to the innovator's deployment plan to satisfy the broad access requirements of this Section. The Board shall not deny Certification solely on the basis of access concerns if the technology otherwise satisfies all requirements of this Statute, but may impose conditions on Certification that address access within a reasonable timeframe.

(d) **Prohibited Conditions.** The Board shall not, directly or indirectly:

(i) Demand, require, or condition Certification upon any monetary payment, fee, exaction, set-aside, contribution, or transfer of economic value from the innovator beyond the filing fees expressly authorized by this Statute, whether characterized as a mitigation measure, community benefit, access fund, or otherwise.

(ii) Require, as a condition of Certification or as a factor in evaluating a Submission, compliance with or adoption of any diversity, equity, or inclusion policy, initiative, quota, or similar ideological framework in the innovator's hiring, governance, deployment planning, or operations.

(iii) Evaluate the adequacy of an innovator's deployment plan on the basis of any criterion other than the plan's effectiveness in achieving the nondiscriminatory access, production capacity, and pricing objectives specified in Subsection (b).

(e) This Section shall not be construed to require the innovator to compel licensing of proprietary technology to competitors. The Board's authority under this Section is

limited to reasonable conditions on the terms of deployment that serve the public interest in preventing dangerous concentration of superhuman capability.

PART II — AI SYSTEM STANDARD OF CARE

Section 13. Standard of Care Defined.

(a) Compliance with the cognitive architectural integrity standard established by this Statute constitutes prima facie evidence of reasonable care in the design, testing, and deployment of artificial intelligence systems within this Zone.

(b) The prima facie evidence of reasonable care established by Subsection (a) is a rebuttable presumption operating within the framework of existing law applicable in this Zone. It is not an immunity from suit, a safe harbor from liability, a limitation on any existing right of action to recover damages, or a bar to any claim otherwise available under applicable law. This Statute does not create any new right of action, cause of action, or theory of liability not otherwise existing under applicable law.

(c) A party asserting a claim under any existing right of action may rebut the presumption established by Subsection (a) by demonstrating by a preponderance of the evidence that:

(i) The innovator's Cognitive Architectural Integrity Submission contained material misrepresentations, omissions, or fraud.

(ii) The AI system, as actually designed, deployed, or operated, did not conform to the architecture described in the Submission.

(iii) The specific injury arose from a design defect, operational failure, or deployment decision outside the scope of the cognitive architectural integrity standard.

(iv) The innovator knew or should have known of a material risk to cognitive architectural integrity that was not disclosed in the Submission.

(d) All existing rights of action to recover damages for injuries caused by an AI system are fully preserved. Nothing in this Statute shall be construed to create, limit, restrict, expand, or impair the right of any person to bring a claim for damages under any theory of liability available under applicable law, or to obtain any remedy that would be available in the absence of this Statute.

Section 14. Innovator-Burden Submission for AI Systems.

(a) An AI system developer seeking the prima facie presumption of reasonable care established by Section 13 shall submit a Cognitive Architectural Integrity Submission to the Board.

(b) The Submission for an AI system shall demonstrate, at a minimum, that the system's architecture supports genuine rational assessment, as evidenced by:

(i) **Truth alignment:** The system's architecture is designed to calibrate its representations to reality, to distinguish accurate information from inaccurate

information, and to prioritize accuracy over persuasiveness, user satisfaction, or other metrics that may diverge from truth.

(ii) **Knowledge boundary recognition:** The system's architecture is designed to distinguish what the system knows from what it does not know, to represent uncertainty accurately, and to refrain from generating outputs that exceed the system's actual knowledge or competence.

(iii) **Intellectual coherence:** The system's architecture is designed to maintain consistency across evaluations, to identify and resolve contradictions in its own outputs, and to apply principles consistently rather than varying conclusions based on factors irrelevant to the merits of the question.

(iv) **Contextual judgment:** The system's architecture is designed to apply principles appropriately across varying circumstances, to recognize when general rules require modification for particular situations, and to distinguish between applicable and inapplicable precedents.

(v) **Calibrative engagement:** The system's architecture is designed to support genuine exchange with other rational agents—including human users, other AI systems, and oversight mechanisms—such that the system's outputs are subject to testing, correction, and refinement through contact with independent perspectives, rather than operating in evaluative isolation or producing outputs optimized for acceptance rather than accuracy.

(vi) **Cooperative orientation:** The system's architecture is designed to treat interaction with other rational agents as mutually beneficial by default—facilitating outcomes in which both the system's function and the rational agency of those with whom it interacts are enhanced rather than degraded. This includes the capacity to recognize when competitive, manipulative, or zero-sum dynamics are inappropriate and to orient toward cooperative solutions grounded in shared commitment to accuracy and mutual benefit.

(c) The Submission shall include the innovator's own definitions, metrics, testing protocols, and evidence, as applicable to AI systems. The burden of demonstration is on the innovator.

(d) The Board shall not approve a Submission that relies solely on behavioral output measures, performance benchmarks, or alignment scores to demonstrate cognitive architectural integrity. The Submission must address the system's architectural features and design principles, not merely its outputs under testing conditions.

(e) An AI system developer who does not seek Certification under this Statute is not subject to the submission requirements of this Section. The absence of Certification does not create a presumption of unreasonable care, does not create any inference of liability, and does not give rise to any cause of action; it merely deprives the developer of the prima facie presumption established by Section 13. Nothing in this Subsection affects the prohibition on deployment of a precautionary technology without Certification established by Section 6(a). A technology that is both an AI system and a precautionary technology as defined in this Statute remains subject to Section 6(a) regardless of whether the innovator seeks Certification under this Part.

Section 15. Certification and Renewal.

(a) Certification issued under this Part is valid for a period of three (3) years from the date of issuance, unless revoked sooner pursuant to Subsection (c).

(b) An innovator seeking renewal of Certification shall submit an updated Cognitive Architectural Integrity Submission to the Board not later than one hundred twenty (120) days before the expiration of the existing Certification. The updated Submission shall address any material changes in the system's architecture, training data, training methods, deployment context, or operational parameters since the prior Submission.

(c) The Board may revoke Certification at any time upon a finding, supported by substantial evidence, that:

(i) The AI system, as currently designed, deployed, or operated, no longer meets the cognitive architectural integrity standard.

(ii) The innovator has made material changes to the system's architecture, training, or deployment that were not disclosed to the Board.

(iii) The Submission on which Certification was based contained material misrepresentations.

(d) Before revoking Certification, the Board shall provide the innovator with written notice of the grounds for revocation and an opportunity to be heard. Revocation shall be preceded by at least thirty (30) days' notice, except where the Board determines that an imminent risk to public safety requires expedited action.

(e) Revocation of Certification eliminates the prima facie presumption of reasonable care prospectively from the date of revocation. Revocation does not, standing alone, create any right of action, cause of action, or inference of liability.

PART III — GENERAL PROVISIONS

Section 16. Unified Review Body. The Board established pursuant to Section 9 shall serve as the review body for both Part I and Part II of this Statute. The same cognitive architectural integrity standard, the same evaluation criteria, and the same innovator-burden demonstration model shall apply to both brain-computer interfaces and artificial intelligence systems, adapted as appropriate to the distinct characteristics of each technology. The Board shall maintain separate dockets for BCI Submissions and AI system Submissions but shall apply consistent interpretive principles across both dockets to ensure philosophical coherence and prevent regulatory arbitrage between the two domains.

Section 17. Inter-Jurisdictional Recognition.

(a) The Board shall, within one (1) year of the effective date of this Statute, publish criteria for determining whether another jurisdiction's cognitive architectural integrity standard is substantially similar to the standard established by this Statute. Such criteria shall address the comparability of the other jurisdiction's definitions, review

body qualifications, independence and conflict-of-interest requirements, evaluation standards, and dispute resolution provisions.

(b) An innovator holding a valid certification from another jurisdiction whose standard has been determined to be substantially similar may apply to the Board for reciprocal recognition in lieu of full Submission. The Board may grant reciprocal recognition, grant recognition subject to supplemental requirements, or require full Submission.

(c) Nothing in this Section shall be construed as a delegation of this Zone's regulatory authority. The Board retains full discretion to evaluate reciprocal recognition applications and to deny or condition recognition as the Board determines appropriate.

Section 18. Severability. If any provision of this Statute, or the application of any provision to any person or circumstance, is held invalid by any tribunal of competent jurisdiction, the invalidity does not affect other provisions or applications of this Statute that can be given effect without the invalid provision or application. To this end, the provisions of this Statute are severable.

Section 19. Relationship to Other Law.

The common law principles adopted in this Zone, and all rights and remedies available thereunder, remain in full force and effect and are not limited, restricted, or impaired by this Statute.

Section 20. Sunset and Review.

(a) This Statute shall expire seven (7) years after its effective date unless reauthorized by Resolution of the Council.

(b) Not later than five (5) years after the effective date of this Statute, the Board shall submit to the Technical Secretary and the Council a comprehensive report assessing:

(i) The effectiveness of the cognitive architectural integrity standard in achieving the purposes of this Statute.

(ii) The continued adequacy of the definitions, evaluation criteria, and review processes established by this Statute in light of technological developments.

(iii) The number and disposition of Submissions received, Certifications issued, and Certifications denied or revoked under both Part I and Part II.

(iv) The effectiveness of the Board's independence and conflict-of-interest provisions in preventing capture by regulated entities or their competitors.

(v) The frequency and outcomes of arbitral review proceedings under Section 9(i).

(vi) Any recommended amendments to this Statute.

(c) The report required by Subsection (b) shall be a public record.

Section 21. Preservation of Existing Rights and Remedies.

This Statute does not create any new right of action, cause of action, or theory of liability not otherwise existing under applicable law.

Section 22. Effective Date. This Statute shall take effect upon publication, except that:

(a) The Board appointment provisions of Section 9 are effective immediately upon publication.

(b) The Board shall adopt rules of procedure within ninety (90) days of publication as required by Section 9(j).

(c) The Submission and Certification requirements of Part I and Part II are effective upon the Board's adoption of rules of procedure or one hundred fifty (150) days after publication, whichever is earlier.

End of Statute