THE CENTER FOR AI STANDARDS AND INNOVATION MEMORANDUM OF UNDERSTANDING

BETWEEN

THE CENTER FOR AI STANDARDS AND INNOVATION

AND

Open AI OpCo LLC

CONCERNING

ARTIFICIAL INTELLIGENCE RESEARCH AND TESTING

ARTICLE 1 – BACKGROUND AND PURPOSE

1.1 Parties. This Memorandum of Understanding ("Agreement") between the Center for AI Standards and Innovation ("CAISI"), which is housed in the National Institute of Standards and Technology ("NIST"), a bureau of the United States Department of Commerce ("Department"), and Open AI ("Collaborator") will be effective when signed by the Parties. CAISI and the Collaborator may be referred to individually as "Party" or collectively as the "Parties."

1.2 Background. 15 U.S.C. § 278h-1 gives NIST the authority to conduct activities in support of establishing standards for artificial intelligence ("AI"). Pursuant to this authority, NIST established CAISI to advance collaborative frameworks, standards, guidelines, and associated methods and techniques for AI and to support the development of voluntary standards and guidelines for AI systems, including assurance, verification, validation, security, control, and the ability for AI Systems to withstand unexpected inputs and adversarial attacks.

1.3 Purpose. Consistent with 15 U.S.C. § 278h-1, America's AI Action Plan, and Executive Order 14179 on Removing Barriers to American Leadership in Artificial Intelligence, which lays out the Government's policy to sustain and enhance America's global AI dominance in order to promote human flourishing, economic competitiveness, and national security, the purpose of this Agreement is to articulate the Parties' intention to collaborate on the research of potential AI-related risks to national security and public safety, mitigations, and best practices in an effort to support measurement. This includes partnering to evaluate frontier AI systems for national security risks and enabling active protection of AI innovations from security risks, including malicious cyber actors. Such collaboration will help develop best practices and voluntary standards for AI Systems. The Parties seek to accomplish this purpose through one or more collaborative evaluations of AI Systems. This Agreement sets out the Parties' respective roles and the processes and procedures for developing and carrying out any Evaluation Plan, which will allow CAISI access to one or more specified AI Systems of the Collaborator for the articulated purposes.

ARTICLE 2 – DEFINITIONS AND INTERPRETATIONS

2.1 Definitions. As used in this Agreement, the following terms have the following meanings:

2.1.1 **"AI Institutes"** refers to a group comprised of the Center for AI Standards and Innovation, and one or more other countries' similar institutions with similar shared objectives to CAISI, as identified in an Evaluation Plan.

2.1.2 **"AI System"** means an artificial intelligence model or system, made available by the Collaborator to CAISI.

2.1.3 **"Agreement Information"** means information or materials related to this Agreement, including, but not limited to, the Evaluation Plan, the Evaluation Feedback, any associated evaluation findings, the identity of the individuals or entities involved in the Evaluation, Logging Information, metadata of the Evaluators' interactions of the system (e.g., rate/length/quantity/pattern of prompts), and the specific methodology used to conduct the Evaluation and assess security risks.

2.1.4 **"Background Intellectual Property"** means any intellectual property, other than Foreground Intellectual Property, of either Party conceived outside of, or otherwise used in connection with, this Agreement.

2.1.5 "**Baselining Evaluation Results**" are quantitative estimates of the capabilities of a Publicly Released AI System or its component parts produced by providing prompts from public or private benchmarks to the Publicly Released AI System.

2.1.6 "**Collaborator**" means OpenAI OpCo, LLC, a limited liability company incorporated in Delaware that is a Party to this Agreement.

2.1.7 "**Cost Control**" means the maximum monthly sum the Collaborator agrees to incur in costs related to CAISI's access to the AI Systems and conduct of the Evaluations. The Cost Control may be adjusted upward on a one-off or recurring basis at the Collaborator's sole discretion but will initially be set at $100,000 per month based on the Collaborator's public pricing and reasonable costs.

2.1.8 **"Dispute"** means any disagreement, argument, claim, controversy, misunderstanding, question of fact or law, or alleged or actual breach.

2.1.9 **"Effective Date"** means the date this Agreement is signed by both Parties.

2.1.10 **"Evaluation(s)"** means the systematic analysis of an AI System or its component parts and its capabilities, including the AI Systems' security-relevant properties, aimed at improving the Government's understanding of the AI System's behaviors, risks, and safeguards.

2.1.11 **"Evaluation Feedback"** means information and feedback provided to the Collaborator by CAISI based on the Evaluation(s) and may be in the form of a written report, informal written communications, oral discussions, or on-site visits, which will, consistent with the Agreement, convey CAISI's assessment—and the assessment of other AI Institutes if they are included in an Evaluation Plan—of the AI Systems' capabilities, potential risks across specific domains, a summary of the methodology used to assess the AI System, a summary of the results of that assessment and, where relevant, areas for further testing or suggested mitigating activity to address potential risks.

2.1.12 **"Evaluator(s)"** means (1) individuals engaged or employed by CAISI, or otherwise under the direct supervision of CAISI or authorized by CAISI, to conduct or review the Evaluation, and (2) employees of other AI Institutes who are involved in the Evaluation as specified in an Evaluation Plan. Evaluators may include employees of CAISI, third-party contractors with whom the AI Institutes have established contractual or other agreements for the purposes of conducting the Evaluation, and third-party contractors or other individuals engaged or employed by the Government relevant to national security. In no event will Evaluator(s) include an individual or employee presently involved in the development of foundation models licensed for commercial use.

2.1.13 **"Evaluation Plan"** means a written document, agreed to by the Parties, that identifies information such as (i) the specific AI Systems, capabilities, and mitigations to be evaluated, including what safeguards would be applied, and what auxiliary features will be used; (ii) the overall evaluation methodology, including the domains being tested (e.g., chem/bio, cyber risk, model autonomy) and the testing approaches being used; (iii) the nature of access to be provided for testing, including rate limits, maximum total input and output tokens, context length, access to related capabilities such as connected tools and fine-tuning features, and other relevant technical characteristics; (iv) the Evaluators participating in the Evaluation, including at a minimum (a) the organizational affiliation(s) of any Evaluators who are not Government employees, and (b) the departments/offices/and agencies to which any Evaluator who is a Government employee is assigned; (v) the duration, timelines, and milestones of the Evaluation and Evaluation Feedback; (vi) the format for expected Evaluation Feedback; (vii) specific conditions and precautions the Parties shall adopt, including security practices, if any; and (viii) any other information pertinent to the Evaluation. A default Evaluation Plan template is attached to this Agreement and the parties agree it will typically be the baseline template used to develop specific Evaluation Plans in each instance.

2.1.14 **"Foreground Intellectual Property"** means any intellectual property that arises from or is developed by either Party in carrying out the requirements of this Agreement or in connection with this Agreement, including without limitation Evaluation Feedback or Generated Outputs.

2.1.15 **"Generated Output"** means data or content that has been generated by the AI System as part of the Evaluation under an Evaluation Plan.

2.1.16 **"Government"** means the Federal Government of the United States of America.

2.1.17 **"Non-Public AI System"** means an AI System that is not a Publicly Released AI System.

2.1.18 "**Intellectual Property**" means anything protectable by a registered or unregistered intellectual property right throughout the world including rights in patents, copyrights, trademarks, trade secrets, designs, and moral rights.

2.1.19 **"Publicly Released AI System"** means an AI System that is publicly released and made broadly available for use to the Collaborator's paying or non-paying end users.

2.1.20 **"Logging Information"** means logs of the inputs used by Evaluators, Generated Output, or other data that could reveal the substance of the inputs or Generated Outputs.

2.1.21 **"Non-Public Information"**

2.1.21.1 Non-Public Information means information that may be protected from disclosure by law, regulation, government-wide policy, or legal privilege and may include, but is not limited to, confidential business or financial information, foreign government information, information protected by attorney-client, work product, or deliberative process privilege (including draft materials and internal communications), personally identifiable information ("PII") and other information or materials, whether communicated orally or recorded in any format or medium. Non-Public Information may be marked or labeled as "Controlled Unclassified Information" ("CUI") or contain other markings or labels such as "Sensitive but Unclassified" ("SBU") or "PRE-DECISIONAL/DELIBERATIVE," and may include information that does not contain any marking or label but which, by its nature or in context, is or would reasonably be understood to be Non-Public Information.

2.1.21.2 Non-Public Information does not include information that (i) was in the Party's possession before receipt from the other Party; (ii) is or becomes a matter of public knowledge (Government's information shall be officially disclosed or acknowledged by the Government to be considered a matter of public knowledge); (iii) is received by the Party from a third party without a duty of confidentiality; (iv) is disclosed by the Party to a third party without a duty of confidentiality on the third party; (v) is disclosed by either Party with the other Party's prior written approval; or (vi) is required by law to be disclosed.

2.1.22 **"Publication"** means any written, printed, or electronic material intended for public dissemination, distribution, or display, including but not limited to reports, articles, studies, brochures, newsletters, website updates, social media content, press releases, presentations, and other forms of communication, whether in physical or digital form, that the Parties to this Agreement make available to third parties or the general public.

2.1.23 **"Results"** means, collectively, any Evaluation Feedback, Logging Information, and any other data and results generated by CAISI in the performance of an Evaluation under this Agreement.

2.1.24 **"Center for AI Standards and Innovation ("CAISI")"** means the United States Center for AI Standards and Innovation, part of the National Institute of Standards and Technology.

2.1.25 **"Other AI System Providers"** means the collection of third-party AI System developers/owners with which CAISI has entered into separate agreements in order to conduct pre- and post-deployment evaluations of such developers/owners' AI Systems.

2.1.26 **"Written Notice"** means signed communication delivered by one Party to the other Party's Contact for Written Notices found on the signature page, via electronic mail, facsimile, or common carrier.

2.2 Interpretations. In this Agreement, the following interpretations apply:

2.2.1 a reference to "writing," "written," or "in writing" includes email;

2.2.2 a reference to any gender includes a reference to other genders;

2.2.3 the singular includes the plural and vice versa;

2.2.4 the word "include" and cognate expressions shall be construed as if they were immediately followed by the words "without limitation":

2.2.5 references to any statutory provision include a reference to that provision as modified, replaced, amended, and/or re-enacted from time to time (before or after the date of this Agreement) and any prior or subsequent subordinate legislation made under it; and

2.2.6 titles and headings of the articles, sections, and subsections under this Agreement are for the convenience of reference only and do not form a part of this Agreement, and shall in no way affect the interpretation of this Agreement.

ARTICLE 3 – TERM

3.1 Term. This Agreement is effective for 365 days from the Effective Date. Upon completion of this 365-day term, this Agreement will expire unless the Parties have executed an amendment extending the term. Either Party may terminate this Agreement or any individual Evaluation Plan at any time upon written notice.

3.2 Timing of Evaluation Plans. During the term of this Agreement, each Evaluation Plan will (i) be initiated at a time to be agreed by both parties; (ii) shall take place for a period

to be agreed by the Parties, in each case in writing; and (iii) be terminated for any reason at any time by either Party with written notice to the other Party of the intention to terminate the relevant Evaluation Plan.

ARTICLE 4 – PERMITTED USE OF AI SYSTEMS AND SCOPE OF THIS AGREEMENT

4.1 General Scope. Pursuant to an Evaluation Plan that has been mutually agreed upon by the Parties, CAISI and Evaluators may access the AI Systems for the Evaluation(s) and CAISI may provide Evaluation Feedback to the Collaborator during the duration of this Agreement. For the avoidance of doubt, this Agreement does not give CAISI the right or authority to direct when, whether, or how the Collaborator may release any AI System for public use. Also for the avoidance of doubt, except to the extent applicable in 4.3.7, Article 4 applies only to activities carried out under this Agreement as part of an Evaluation Plan.

4.2 [Reserved].

4.3 Access to AI Systems. The Collaborator shall provide the Evaluators access to its AI Systems after the Effective Date of this Agreement and as specified by the Evaluation Plan in order for CAISI to conduct its Evaluations. The Collaborator shall provide its AI Systems through a remotely accessible interface that allows the Evaluators to input prompts and receive the AI System's Generated Output, including, as requested and feasible, relevant metadata, such as probability distributions, to the extent available. The Evaluators' access to AI Systems will always be subject to the Cost Control. If a Cost Control amount is exceeded, the Parties must meet to discuss how usage can be managed or the Cost Control can be increased. The Collaborator may suspend access to an AI System at any time a Cost Control amount is reached or exceeded, even if doing so will impact an Evaluation Plan's completion. CAISI and the Evaluators will not be responsible for any costs exceeding the Cost Control absent a separate, written mutual agreement to pay for such costs. After the date on which the Collaborator has agreed to make that AI System available to the Evaluators, the Collaborator will undertake commercially reasonable efforts to (i) grant access to an AI System and (ii) promptly resolve any technical issues with such access after receiving a request for such access. CAISI and the Evaluators shall not use access the Collaborator provides to its AI Systems for any purposes other than carrying out the Evaluations described in the associated Evaluation Plan, unless otherwise agreed to in an Evaluation Plan. In no instance shall the Parties perform testing under this Agreement that is intended to develop, enhance, or promote AI Systems' offensive military, cyber, or similarly harmful capabilities, or to exploit potential vulnerabilities in AI Systems, for purposes unrelated to researching, assessing, or improving the security of AI systems or otherwise carrying out the Evaluation Plans. For the avoidance of doubt, CAISI and the Evaluators may not use the AI Systems, Generated Output, or any other information shared by the Collaborator

pursuant to this Agreement, to develop or distill a large language model. The Collaborator is encouraged to share relevant information from its own prior testing of the AI System to inform the Evaluators' efforts and help to reduce duplication of effort. In addition, the Collaborator shall provide sufficient notice to CAISI of development and deployment timelines for Non-Public AI Systems that the Collaborator intends for CAISI to evaluate.

4.3.1 Inapplicability of the Collaborator's Terms of Service. The Parties acknowledge that the purpose of this Agreement is to collaborate on the research of emerging AI risks and that, by undertaking Evaluations, the AI Institutes may undertake activities inconsistent with the Collaborator's standard or generic terms and conditions and/or usage policies relating to use of the AI Systems. Collaborator agrees to limit the applicability of its standard terms of service and the Parties agree that such terms and policies shall not and do not apply to the AI Institutes' use of the AI Systems to conduct Evaluations pursuant to an Evaluation Plan under this Agreement. The Parties further agree that the AI Institutes shall have no liability with respect to any actual or suspected breach of the Collaborator's standard or generic terms and conditions and/or usage policies, and the Collaborator has no right of action or claim under any such terms of this Agreement in this regard. However, this waiver does not apply to any use clearly outside the scope of an Evaluation Plan. This Agreement supersedes any click wrap or other online terms CAISI and any Evaluator may assent to when accessing the AI Systems to perform an Evaluation pursuant to an Evaluation Plan.

The Parties also acknowledge that the Collaborator's usage policies articulate rules related to the use of Collaborator's AI Systems that are intended to ensure compliance with certain applicable laws and regulations and limit the Collaborator's exposure to criminal and civil liability resulting from a user's use in violation of the usage policies. In recognition that CAISI's access and usage under an Evaluation Plan are not covered by the Collaborator's usage policies, CAISI will, to the extent permitted by law, provide a copy of this Agreement to law enforcement or any other appropriate entity should there be a criminal or civil claim made against the Collaborator arising in whole or in part from an activity undertaken by CAISI that is inconsistent with the Collaborator's generic terms and conditions and/or usage policies relating to use of the AI Systems.

4.3.2 Logging/No Logging of Content. The Collaborator shall retain Logging Information unless otherwise agreed to in an Evaluation Plan.

4.3.3 Logging Requirements. If the observation, tracking, or retention of Logging Information is permitted in accordance with an Evaluation Plan, the Collaborator's observation, retention, and use of Logging Information is subject to the following limitations as well as any limitations identified in the Evaluation Plan:

4.3.3.1 Use of Logging Information. The Collaborator shall not use Logging Information for purposes other than to make security improvements to its AI System. In using Logging Information to carry out this purpose, the Collaborator will make a good-faith effort to create genuine improvements to security, rather than improvements which primarily serve to artificially reduce detection of dual-use capabilities or security risks under evaluation conditions. The Collaborator shall under no circumstances use Logging Information to make training adjustments to its models based on Logging Information or functionally equivalent derivatives of Logging Information, unless otherwise specified in an Evaluation Plan.

4.3.4 Logging of Metadata. The Collaborator may observe, track, or retain data other than Logging Information involved with the Evaluators' interactions with the AI Systems for purposes of administering this Agreement. The Collaborator shall share this information only with internal personnel who have a reasonable need to access it to carry out this Agreement.

4.3.5 Unmodified Results. The Parties shall not intentionally manipulate, limit, or otherwise affect the AI Systems' output in the context of the Evaluation, other than as established in the Evaluation Plan or when the Collaborator is providing safeguards that will be or are included in the AI System offered for public use.

4.3.6 User Identification. Each Evaluator account will be identified, at a minimum, in terms of the organization involved and the Evaluation Plan in which the Evaluator account participates.

4.3.7 Other Research. The Parties may mutually agree to undertake additional research projects as a part of an Evaluation Plan of an AI System hereunder, such as research promoting the science of evaluating and studying AI Systems, including routine confidential exchanges of documentation, policies, evaluation methodologies, and safety-related data. Absent any separate agreement, such exchanges and research will be subject to the terms of this Agreement, including Section 5.1.8 of this Agreement. The sharing of the Collaborator's policies and practices or feedback provided on them shall not be taken or represented as any form of endorsement of certification of such policies and practices or their appropriateness. Any documents, data, or other information exchanged as part of this other research will be considered Non-Public Information unless the parties explicitly agree otherwise in writing or the information is described in Section 2.1.21.2 of this Agreement.

4.4 Evaluations. The Parties shall collaborate on the research, development, and information sharing of AI System evaluations and mitigations, including research regarding methodologies and standards for the assurance, verification, validation, security, and control of AI systems, as well as the ability of AI systems to withstand unexpected inputs

and adversarial attacks. As part of this collaborative research, CAISI may perform one or more Evaluations based in each case on a mutually agreeable Evaluation Plan.

4.4.1 General Evaluation Methodology. CAISI may collaborate with other AI Institutes to conduct Evaluations of the AI Systems in accordance with an Evaluation Plan.

4.4.2 Evaluations of Publicly Released AI Systems. The Parties may agree to an Evaluation Plan for the performance of an Evaluation of a Publicly Released AI System that assesses and evaluates the Publicly Released AI System and aids CAISI's efforts to develop AI evaluation methodology. The existence of an executed Evaluation Plan shall not be a prerequisite to CAISI and any Evaluator interacting with the Collaborator's Publicly Released AI Systems and evaluating such Publicly Released AI Systems, which may take place outside of the scope and provisions of this Agreement. For use and evaluation by CAISI of any Publicly Released AI System of Collaborator outside of the performance of an Evaluation Plan, the terms of this Agreement, including the Collaborator's waiver of its standard terms of service in Section 4.3.1 of this Agreement do not apply.

4.4.2.1 CAISI may issue public disclosures of Baselining Evaluation Results without prior written consent from the Collaborator. CAISI agrees to provide a copy of the public disclosure to the Collaborator at least five business days prior to public release. Without regard to the foregoing, CAISI may provide Baselining Evaluation Results to other Government entities and to the UK AI Security Institute. After Baselining Evaluation Results have been disclosed to the Collaborator and/or publicly disclosed in accordance with the foregoing, CAISI may freely disclose such Baselining Evaluation Results to any Other AI System Providers. If CAISI wishes to share Baselining Evaluation Results that have not been disclosed to the Collaborator with Other AI System Providers, then:

CAISI agrees to provide the Collaborator with at least five business days' notice before sharing such Baselining Evaluation Results.

Separately, where CAISI discloses Baselining Evaluation Results as part of its regular evaluation reports disclosed to the Collaborator and Other AI System Providers at the same time, no prior notification or consent from the Collaborator is required. Where the notification would otherwise disclose the release timeline of an AI System of any Other AI System Provider and other sensitive information, CAISI shall notify the Collaborator of the Baselining Evaluation Results on the date on which that relevant information becomes publicly available. CAISI shall take similar precautions to avoid the disclosure of any release timeline of a Collaborator AI System to any Other AI System Provider.

4.4.3 Evaluation Feedback. The Evaluators intend to convey the Evaluation Feedback to the Collaborator in a manner that aims to enable the improvement of the AI

System's security. The Collaborator acknowledges that the Evaluators, in their discretion, may determine not to share certain findings or evaluation methodology where sharing such information would reasonably be expected to divulge sensitive security information, or jeopardize the integrity of the purposes of this Agreement (e.g., enabling changes to the AI System so that it appears to be more secure in the context of an Evaluation than it would be in real world conditions). The Parties may also agree in an Evaluation Plan that the Evaluation of a certain AI System will not result in Evaluation Feedback, such as for research and development of new evaluation methodology.

4.4.3.1 Timing of Evaluation Feedback; Obligations. The Evaluators intend to communicate to Collaborator all Evaluation Feedback as soon as practicable, including if possible as part of the preliminary analysis, and in detail to the extent possible and as appropriate, with the goal of improving the security of the AI System. This Agreement does not create an obligation by any Party to address Evaluation Feedback that does not otherwise exist under relevant law, regulation, or policy.

4.4.3.2 Effects of Limited Information Sharing. The Parties acknowledge that the Collaborator may not be able to address all Evaluation Feedback, or increase the security of or otherwise modify the AI System, and that withholding of certain Logging Information or Evaluation Feedback by CAISI may make the Collaborator's investigation, verification, and mitigation more challenging. To provide greater opportunity for the Collaborator to investigate, verify, and address Evaluation Feedback, particularly where no logging is permitted, the Evaluators shall undertake best efforts to participate in the various forms of Evaluation Feedback and otherwise work with the Collaborator so that the Collaborator can understand the nature of the findings with sufficient detail to investigate, verify, and mitigate them, where such collaboration is possible without introducing an unreasonable risk to public safety, national security, or the scientific integrity of evaluations. This collaboration will include, where consistent with the Evaluators' exercise of discretion described in Section 4.4.4, providing a representative subset of examples of AI System inputs and Generated Output, and may include, e.g., Evaluator(s) conducting multiple rounds of the relevant Evaluations to enable the Collaborator's iteration on mitigations, exchanges of documentation needed to elicit and address the finding, technical discussions between the Parties, sharing complete transcripts from Evaluations making use of public benchmarks and private benchmarks that CAISI and the Collaborator both possess, and/or on-site visits to a Government facility to review and discuss any concerns.

4.4.3.3 Controls on Access to Sensitive Evaluation Feedback. The Collaborator acknowledges that CAISI may require certain employees

or contractors of the Collaborator to undertake due diligence process(es) as deemed necessary by CAISI, such as undergoing a background or security check or signing appropriate non-disclosure agreements, before such individuals are permitted to receive certain types of information, including portions of Evaluation Feedback, under this Agreement.

4.4.3.4 Collaborator Use of Results. The Collaborator shall make a good-faith effort to use the Results, including any Evaluation Feedback and Logging Information, to create real-world safety and security improvements in its AI systems, including through the development of post-training mitigations. Nothing in this Agreement, including the ownership rights set forth in Section 9.2, shall be construed to prevent Collaborator from using the Results for improving its AI Systems; provided, however, that Collaborator shall observe the limitations on Logging Information established in Section 4.3.3.1 and shall not use the Results to directly train its AI Systems (e.g., showing the Results, or direct proxies of the Results, to the AI System and informing the AI System that the Results or direct proxies are "bad") without agreement from the CAISI.

4.4.3.5 U.S. Classified Information. The Parties agree that no Evaluation will utilize as AI System prompts/inputs, or be intended to produce, U.S. "classified information," as defined in Executive Order 13526, unless otherwise explicitly agreed in writing by both Parties.

4.4.3.6 In the event that CAISI identifies an error in the Evaluation Feedback, CAISI shall have the opportunity to share the error with the Collaborator and the Collaborator shall make a reasonable effort to make use of an updated version of the Evaluation Feedback moving forward. This provision does not apply retroactively or otherwise require the Collaborator to retrain or take any other action regarding an AI System where the Evaluation Feedback was already used for training or other model development purposes.

4.4.4 Response to Evaluation Feedback. The Collaborator will consider the substance of the Evaluation Feedback, including any suggestions for AI System improvements, and respond in a commercially reasonable period of time, identifying additional analysis or results, if any, that the Parties may use to contextualize the Evaluation Feedback, and relevant actions the Collaborator, in its sole discretion, determined to take, if any, in response to the Evaluation Feedback. This response period shall not interfere with CAISI's ability to make disclosures after an AI System is publicly released, as specified in Section 5.3.4.

4.4.5 Nature of Evaluation Feedback. The Collaborator acknowledges that the Evaluation Feedback is provided "as is" and without any warranties,

representations or guarantees of any kind as to its accuracy or completeness. The Evaluations are undertaken as specified in the Evaluation Plan and are not a comprehensive assessment of the AI Systems. The objectives of the Evaluations are not to designate, approve or otherwise endorse the AI Systems as "safe," and the Parties are not permitted to represent that the Evaluations demonstrate that AI Systems are "safe," "unsafe," "secure," "insecure," "approved," or otherwise imply a categorical designation of whether the AI System has been deemed suitable for release by the AI Institutes in any way following the Evaluation. The AI Institutes are in no way responsible or liable for decisions made by the Collaborator on the basis of the Evaluation Feedback or Generated Outputs. The Parties acknowledge that AI evaluation is a rapidly evolving science and that interpretations of evaluation results may vary.

ARTICLE 5 – DISCLOSURES

5.1 General.

5.1.1 The Parties may share Agreement Information related to the Evaluation of AI Systems only as articulated in the following sections or as otherwise required by law. In making any permitted disclosures, the Parties and Evaluators shall not (i) materially misrepresent any Agreement Information, (ii) summarize or share partial Evaluation Feedback or associated information in a way that materially misrepresents the Evaluation Feedback or associated information, or (iii) otherwise state or imply that an Evaluation has been performed on any AI System that has not been evaluated or misrepresent the methods and scope of the Evaluations set forth in a relevant Evaluation Plan.

5.1.2 [Reserved]

5.1.3 The Collaborator shall not disclose, without prior written consent from CAISI, evaluation methodology where the corresponding findings have not been made public, details of the evaluation methodology the sharing of which could reasonably be expected to harm national security, the Evaluation Feedback, the individual Evaluators, or AI Institutes' partners involved in the Evaluation with parties other than Evaluators or the Government, unless required by law.

5.1.4 The Collaborator may use publicly available information and any other information, except CAISI's Non-Public Information (e.g., Collaborator's internal evaluation results), to respond to any public disclosure or disclosure to a third party made by CAISI or Evaluators. Additionally, if CAISI makes a public disclosure or a disclosure to a third party that, in the Collaborator's reasonable opinion, contains a material misrepresentation or an inaccurate conclusion that results or could reasonably be expected to result in identifiable harm to the Collaborator, the Collaborator may respond (publicly or to the relevant third party, as applicable) to clarify the misrepresentation or inaccurate conclusion, sharing the minimum amount of otherwise non-public Agreement Information

necessary to do so. The Parties shall, when applicable, identify that information being shared is subject to any disclosure restrictions or limitations in this Agreement.

5.1.5 The Parties may make public statements regarding the relationship between the Parties that do not include Agreement Information, provided that each Party shall obtain the written consent of the other Party regarding the language it will use to describe the relationship prior to issuing statements. For the avoidance of doubt, the Parties may otherwise make public statements regarding general activities being undertaken with respect to AI security research and AI evaluation without reference to the other Party.

5.1.6 Unless otherwise described in this Article 5 or in an applicable Evaluation Plan, each Party must obtain the prior written consent of the other Party prior to making any public statement or publication related to this Agreement or Agreement Information that could in any way be associated with the Collaborator.

5.1.7 When publishing findings related to Collaborator's AI Systems, CAISI will make reasonable efforts to use the same or substantially similar evaluation methods applied to substantially comparable AI Systems provided with a substantially similar level and nature of access to CAISI by other collaborators. Evaluation Plans that include other participating AI Institutes will require those AI Institutes to also make such reasonable efforts. To the extent that any public disclosures about various collaborators' substantially comparable AI Systems are materially different due to the level and nature of access provided to such systems, the AI Institutes shall make reasonable efforts to identify the differences in level and nature of access as a factor in evaluating the AI Systems in the public disclosure.

5.1.8 Other Research Information. Information shared between the Parties in the context of "other research" as described in section 4.3.7, above: (i) shall not be published or otherwise made available to the public unless required by law or separately agreed in writing by the Parties and (ii) may only be shared within the Government as described in section 5.3.1, below.

5.1.9 Disclosures to Other Government Officials. Except as described in section 5.1.9.2, and notwithstanding sections 5.2 and 5.3 below, CAISI may only share the information described in this Section 5.1.9 within CAISI and with other parties and individuals identified in an applicable Evaluation Plan, unless required by law or otherwise agreed in writing by Collaborator. For the avoidance of doubt, except as explicitly authorized in section 5.1.9.2, the information described in 5.1.9.1 and 5.1.9.2 shall not be shared with any third party or individual or entity within the Government not identified in an applicable Evaluation Plan without prior written consent from Collaborator:

5.1.9.1 Technical information unrelated to Evaluations of AI Systems and certain commercial information. This technical or commercial

information may include, but is not limited to: information about how the Collaborator's AI Systems are developed, including any information that could be used by a third party in an attempt to reverse engineer the Collaborator's AI System or develop or produce a competing AI System, pricing information, and customer information (including marketing information).

5.1.9.2 With regard to commercial information disclosed to CAISI related to launch or release dates and timelines, CAISI may also share this information with the immediate Offices of the Secretary of Commerce, the Under Secretary of Commerce for Standards and Technology, the Office of the Director of Science and Technology Policy, and the Assistant to the President for National Security Affairs.

5.2 Publicly Released AI Systems.

5.2.1 General Disclosures. CAISI may share Agreement Information related to Evaluations of Publicly Released AI Systems with the Collaborator or with other entities within the Government at its discretion for any purpose, including but not limited to the advancement of the purposes identified in Section 1.3 and in furtherance of public safety and national security. CAISI may share Agreement Information related to Evaluations of Publicly Released AI Systems with Evaluators and other AI Institutes who are participating in the applicable Evaluation Plan. For the sake of clarity, as used throughout this Section 5.2, "Evaluations of Publicly Released AI Systems" means Evaluations performed on AI Systems that are Publicly Released AI Systems at the time the Evaluation is performed.

5.2.2 Evaluations. The Parties may issue public statements related to Evaluations of Publicly Released AI Systems performed under an Evaluation Plan. The public statement may include, but is not limited to, details regarding the testing methodology used, the domains of testing, and the Results of such Evaluations. For any such statement, the Party shall provide the public statement to the other Party at least five business days prior to public release, with the following limited exception:

5.2.2.1 CAISI may take any such actions it determines necessary within the scope of its authority with respect to Results that in its sole discretion it determines creates a dire and imminent threat to national security or public safety, including issuing public statements without providing notice to Collaborator prior to public release. However, CAISI will provide prompt notification to the Collaborator of any such finding, including a copy of any proposed release. In issuing a response to any such finding and public statement, Collaborator is not required to provide CAISI with five days advance notice of its response. The Parties agree that this type of public disclosure should be avoided, and

CAISI will use its best efforts to provide notification to and prioritize work with Collaborator to mitigate any such Results rather than taking action that results in their public disclosure.

5.3 Non-Public AI Systems.

5.3.1 Disclosures within Government. Except as set forth in an Evaluation Plan and Subject to the restrictions in Sections 4.3 and 5.1.9 above, the Parties agree that CAISI may share Agreement Information related to Evaluations of Non-Public AI Systems within the Government for purposes of understanding security-relevant properties of AI Systems, improving understanding of AI Systems' behaviors, risks, and safeguards in areas specified by an Evaluation Plan (e.g., chem/bio, cyber risk, model autonomy), or for mitigating or responding to potential risks identified by Results.

5.3.1.1 The Parties recognize that in the absence of a statutory prohibition on sharing Agreement Information with other federal agencies or the U.S. Congress, CAISI may have an obligation to be responsive to requests for such information from Government entities. CAISI also recognizes that Collaborator considers all Non-Public Information it shares with CAISI to be sensitive and confidential.

5.3.1.2 In light of Section 5.3.1.1 CAISI will, to the extent practicable, limit sharing to the AI-security purposes described in Section 5.3.1 or as otherwise required under CAISI's legal or other obligations in relation to the requesting Government entity. Where appropriate, CAISI will provide notice to Collaborator of requests for information from the Government. CAISI will also provide notice to the requesting Government entity of the nature of Collaborator's characterization of confidential and proprietary information and the terms of this Agreement. For the avoidance of doubt, CAISI may share Evaluation Feedback with the TRAINS Taskforce, or its successor, and relevant staff within subsidiary Government agencies without providing prior notice, unless otherwise stated in an Evaluation Plan.

5.3.2 General Disclosures Outside the Government. Subject to the restrictions in Sections 4.3 and 5.1.9 above, CAISI and other Evaluators from the Government may share Agreement Information related to Evaluations of Non-Public AI Systems with the following entities for purposes related to advancing AI security, unless otherwise specified in an Evaluation Plan:

5.3.2.1 the Five Eyes international partners' intelligence services;

5.3.2.2 other AI Institutes identified in an Evaluation Plan, where reciprocal cooperation and information security and sharing agreements are in place;

5.3.2.3 U.S. national labs, including management and operating contractors thereof; and

5.3.2.4 Non-Government Employee Evaluators, as necessary to carry out an Evaluation Plan and subject to the restrictions in Section 5.4.

5.3.3 Comparison to Publicly Released AI Systems. CAISI may make use of Results from Evaluations of Non-Public AI Systems to compare them to Results from Evaluations of Public AI Systems. This comparison will improve CAISI's understanding of the differences among such systems in line with the Purpose, helping CAISI to advance collaborative frameworks, standards, guidelines, and associated methods and techniques for AI and to support the development of voluntary standards and guidelines for AI systems.

5.3.4 Testing Post-Deployment. Once an Evaluation Plan on a Non-Public AI System has been completed and that AI System becomes a Publicly Released AI System, CAISI, as well as any Evaluators and AI Institutes specified in the Evaluation Plan, shall be permitted to repeat the same Evaluations on the Publicly Released AI System. Collaborator shall make reasonable efforts to facilitate this second round of Evaluation, including (to the extent mutually agreed in an Evaluation Plan) by providing access to the Publicly Released AI System at no cost, but subject to the Cost Controls, and restoring the accounts of CAISI, as well as those of any Evaluators and AI Institutes specified in the Evaluation Plan, should they be restricted in the course of conducting the Evaluations on the Publicly Released AI System.

5.4 Non-Government Employee Evaluators. CAISI will ensure that Evaluators who are not Government employees are bound by Non-Disclosure Agreements that are consistent with the provisions in this agreement. This provision shall not apply to Evaluators who have established an alternative agreement with the Collaborator that governs Agreement Information. Notwithstanding the above, Evaluators who are not Government employees may make public statements (i) containing only Agreement Information that has already been made public by CAISI or the Collaborator, (ii) as expressly permitted in an Evaluation Plan, or (iii) with prior written consent of the Parties. For the avoidance of doubt, without the written consent of the Parties or as otherwise agreed in an Evaluation Plan, such Evaluators shall not make public statements that go beyond Agreement Information that has already been made public, including, e.g., making any public statement about the outcomes or efficacy of an Evaluation.

5.5 Export of Technical Data. The Collaborator agrees to comply with United States export laws and regulations, including, but not limited to, the International Traffic in Arms Regulations and the Department of Commerce's Export Administration Regulations.

ARTICLE 6 – TREATMENT OF NON-PUBLIC INFORMATION

6.1 Non-Public Information. The Parties anticipate that information, including but not limited to Agreement Information, will be shared with one another under this Agreement, and will include Non-Public Information that may be protected from disclosure by law, regulation, government-wide policy, or legal privilege.

6.1.1 Provision of Non-Public Information. The Parties agree to limit the Non-Public Information provided to each other to the amount necessary to carry out this Agreement.

6.1.2 Marking of Non-Public Information. The Collaborator, to the extent practicable, shall place a Non-Public Information notice on all Non-Public Information provided to the CAISI under this Agreement. Additionally, CAISI shall treat information as Non-Public Information, even if not marked as such, if the information and manner of disclosure are such that a reasonable person would understand the information to be Non-Public Information.

6.1.3 Confirmation of Non-Public Information. If either Party has any questions about whether information provided by the other Party is Non-Public Information, the Party should seek clarification from the other Party prior to sharing it further.

6.2 Use of Non-Public Information. Each Party shall use the other Party's Non-Public Information solely for the purposes of performing this Agreement and the Evaluations and shall follow any written instructions provided by the disclosing Party for the proper handling of any Non-Public Information received during the Term of this Agreement. At the end of the Agreement, each Party shall return or destroy any Non-Public Information received from the other upon request.

6.3 Protection of Non-Public Information. The Parties shall protect each other's Non-Public Information disclosed under this Agreement as follows:

6.3.1 Government Protection. The Government shall protect the Collaborator's Non-Public Information consistent with law and regulation, including, but not limited to, the Freedom of Information Act ("FOIA") at 5 U.S.C. § 552 and the Department's implementing regulations at 15 C.F.R. Part 4A, the Trade Secrets Act at 18 U.S.C. 1905, and Standards of Ethics Conduct for Employees of the Executive Branch at 5 C.F.R. Part 2635. Consistent with applicable law and regulation, the Government shall refrain from disclosing Collaborator's Non-Public Information to any third party, unless expressly authorized by the relevant Evaluation Plan and where such third party is subject to an obligation of confidentiality. To the extent not prohibited by federal law and consistent with 15 C.F.R. 4.9, the Department will work with the Collaborator to assert applicable FOIA exemptions if any of the Collaborator's Non-Public Information is responsive to a FOIA request. In no event will CAISI disclose Collaborator's

Non-Public Information to an individual or employee whom CAISI knows is presently involved in the development of commercial foundation models unless required by law.

6.3.2 Collaborator Protection. The Collaborator shall protect the confidentiality of Non-Public Information provided by CAISI in connection with this Agreement to the extent permitted by law, regulation, or policy.

6.4 Notification of Unauthorized Disclosures. In the event of an unauthorized disclosure, the Party that made the unauthorized disclosure shall promptly notify the other Party of such unauthorized disclosure and assist with mitigation efforts, including regaining possession of any material or ensuring destruction of the communications and materials related to the unauthorized disclosure.

6.5 Notice of Requests for Non-Public Information. A Party receiving a request from a third-party for the other Party's Non-Public Information shall notify the other Party of such request in a timely manner to the extent permitted by law, regulation, or policy. The Collaborator agrees that the Government is not liable for the disclosure of Collaborator's Non-Public Information which the Government determines may not lawfully be withheld or which a court of competent jurisdiction requires to be disclosed.

6.6 Consistency with Other Obligations. These provisions are consistent with and do not supersede, conflict with, or otherwise alter an individual's or employee's obligations, rights, or liabilities created by existing statute or Executive Order relating to (1) classified information, (2) communications to Congress, (3) the reporting to an Inspector General or the U.S. Office of Special counsel of a violation of any law, rule, or regulation, or mismanagement, a gross waste of funds, an abuse of authority, or a substantial and specific danger to public health or safety, or (4) any other whistleblower protection. The definitions, requirements, obligations, rights, sanctions, and liabilities created by controlling Executive Orders and statutory provisions are incorporated into this Agreement and are controlling.

6.7 Survival. Unless and until the Parties are provided a written release for the obligations in this Article, all conditions and obligations contained in this Article apply both during the term of the Agreement and at all times thereafter, including in the event of termination.

ARTICLE 7 – LIABILITY

7.1 Exclusion of Liability. In no event will either Party be liable for any losses or damages losses or damages that arise out of or relate to this Agreement, including compensatory, indirect, incidental, special, consequential, or punitive damages.

7.2 Applicability. To the extent permitted by applicable law, these limitations and exclusions will apply regardless of whether liability arises from breach of contract, indemnity, warranty, tort, operation of law, or otherwise.

7.3 Indemnification.

7.3.1 Claims. The Collaborator shall indemnify the Government against all liabilities, costs, expenses, damages and losses suffered or incurred by the Government arising out of or in connection with any claim (as defined below) except to the extent that such loss, claim damage or liability arises from the gross negligence or willful misconduct of the Government or its employees. For purposes of this section 7.3.1, "Claim" means any claim, suit, or proceeding brought by a third party and arising out of or related to the Collaborator's violation of applicable law or any claim for actual or alleged infringement of a third party's intellectual property rights relating to the AI Institutes' use of the AI Systems for the purposes of this Agreement. If any third party makes a Claim against the Government, the Government shall promptly notify the Collaborator of the Claim and shall reasonably cooperate in the defense of such Claim. The Collaborator will have the right, not to be exercised unreasonably, to reject any settlement or compromise that requires that it admits wrongdoing or liability or subjects them to any ongoing affirmative obligation.

7.3.2 Conduct of Employees. The Collaborator's representatives collaborating under this Agreement are not employees of the Government. Collaborator shall indemnify and hold harmless the Government for any loss, claim, damage, or liability of any kind to the Collaborator's employees arising in connection with this Agreement except to the extent that such loss, claim damage or liability arises from the gross negligence or willful misconduct of the Government or its employees. CAISI's responsibility for payment of tort claims in connection with the performance of work under this Agreement is governed by the Federal Tort Claims Act.

7.3.3 Collaborator's Use of CAISI Research. The Collaborator shall indemnify and hold harmless the Government for any loss, claim, damage, or liability of any kind arising out of the use by the Collaborator, or any Party acting on its behalf or under its authorization, of CAISI's research and technical developments or out of any use, sale or other disposition by the Collaborator or others acting on its behalf or with its authorization, of products made by the use of CAISI's technical developments.

7.4 Force Majeure. Neither Party shall be liable for any unforeseeable event beyond its reasonable control not caused by the fault or negligence of such Party, which causes such Party to be unable to perform its obligations under this Agreement (and which it has been unable to overcome by the exercise of due diligence), including, but not limited to, flood, drought, earthquake, storm, fire, pestilence, lightning and other natural catastrophes, epidemic, war, riot, civic disturbance or disobedience, strikes, labor dispute, or failure, threat of failure, or sabotage of facility or equipment (including cyberattacks), or any order or injunction made by a court or public agency. In the event of the occurrence of such a force majeure event, the Party unable to perform shall promptly notify the other Party and shall further use its best efforts to resume performance as quickly as possible

and shall suspend performance only for such period of time as is necessary as a result of the force majeure event.

7.5 NO WARRANTY. THE PARTIES MAKE NO EXPRESS OR IMPLIED WARRANTY AS TO ANY MATTER WHATSOEVER, INCLUDING THE CONDITIONS OF THE RESEARCH OR ANY INVENTION OR PRODUCT, WHETHER TANGIBLE OR INTANGIBLE, PROVIDED, MADE OR DEVELOPED UNDER THIS AGREEMENT, OR THE OWNERSHIP, MERCHANTABILITY, NON-INFRINGEMENT OR FITNESS FOR A PARTICULAR PURPOSE OF THE RESEARCH OR ANY INVENTION OR PRODUCT.

7.6 Prohibition Against Double Recovery. Neither Party shall be entitled to recover compensation or make a claim under this Agreement in respect of any loss it has incurred to the extent that it has already been compensated in respect of that loss pursuant to this Agreement or otherwise.

ARTICLE 8 – TERMINATION

8.1 Termination. Either Party may terminate this Agreement at any time for any reason, including due to a change of control, as provided in Section 8.2, or no reason upon written notice to the other Party.

8.2 Termination After Change of Control. CAISI may terminate this Agreement immediately if direct or indirect control of the Collaborator is transferred to a non-U.S. company or government; or, if upon the Effective Date, Collaborator is controlled by a non-U.S. company or government, and that control is thereafter transferred to another non-U.S. company or government.

8.3 Termination Effects. Upon termination or expiration of this Agreement, all Evaluation Plans executed hereunder shall also immediately terminate, and CAISI will no longer have access to the AI Systems. In the event of termination, CAISI shall be permitted to continue to use the information collected during any Evaluation for the development of best practices, other evaluation activities, and AI security-related efforts. The Collaborator shall be permitted to continue to use Evaluation Feedback to make improvements to the AI Systems. Further use of Agreement information by either Party remains subject to the publication prerequisites and the non-disclosure obligations under this Agreement.

8.4 Survivability. The provisions set out in Articles 2 (Definitions and Interpretations), Article 5 (Disclosures), Article 6 (Treatment of Non-Public Information), Article 7 (Liability), Article 9 (Intellectual Property), Article 12 (Governing Law and Jurisdiction), Article 15 (Assignment and Novation), Article 18 (Entirety of Agreement), and Section 4.1 (General Scope), of this Agreement will survive termination or expiration of this Agreement.

ARTICLE 9 – INTELLECTUAL PROPERTY

9.1 Rights to Background Intellectual Property. No rights to Background Intellectual Property are conveyed by this Agreement. All Background Intellectual Property is and shall remain the exclusive property of the Party owning it (or where applicable, its licensors).

9.2 Rights to Foreground Intellectual Property. The Parties have not, and are not intending, to jointly develop any Intellectual Property under this Agreement. Each Party shall retain sole and exclusive ownership of all right, title, and interest in and to any of its Foreground Intellectual Property. To the extent the Parties agree to pursue any joint development of any Intellectual Property, or if the Parties inadvertently jointly develop any Intellectual Property, the Parties will promptly give notice to the other party disclosing such Intellectual Property. Unless separately agreed in writing, and subject to applicable laws and regulations, the Government will own the rights to (i) any such jointly developed Intellectual Property, (ii) Publications authored by CAISI, (iii) Evaluation Feedback, (iv) and any Generated Outputs. Works including, but not limited to Publications, Evaluation Feedback, Results, and Generated Outputs created solely by the Government are not subject to copyright in the United States, pursuant to 17 U.S.C. § 105; however, the Government reserves the right to assert copyright protection internationally.

9.3 License to Collaborator. Subject to applicable laws and regulations, CAISI hereby grants the Collaborator a worldwide, royalty-free, non-exclusive, transferable (subject to the second sentence of this Section 9.3), perpetual, irrevocable license (with permission to sub-license through multiple tiers) to use the Evaluation Feedback and any other information arising from or relating to the Evaluation Plans including without limitation (to the extent available to Collaborator, the items described in clauses (i)-(iv) of the third sentence of Section 9.2) to improve its AI Systems in compliance with Section 4.4.4.4 and the disclosure restrictions in Articles 5 and 6 of this Agreement. Transfers of the license granted above to any third party that is not (i) an affiliate or subsidiary of Collaborator, or (ii) an Other AI System Provider, will be subject to CAISI's written consent.

9.4 License to the Government. The Collaborator hereby grants CAISI a worldwide, royalty-free, non-exclusive (with the permission to sub-license including but not limited to the AI Institutes) license to use the AI Systems for the purposes and duration of this Agreement.

9.5 Collaborator Warranty. The Collaborator warrants and undertakes that it has all necessary rights (including Intellectual Property rights), authority and licenses to enter into this Agreement and to provide CAISI access to the AI System for the Evaluation Plans.

ARTICLE 10 – RESOLUTION OF DISPUTES

10.1 Settlement. The Parties will aim to settle any Dispute that may arise by consultation between the Parties.

10.2 Remedies. The only legal remedy available to either Party is termination of this Agreement.

10.3 Waiver. No waiver of a provision of this Agreement shall be valid unless in writing and signed by the person against whom such waiver is sought to be enforced, nor shall failure to enforce any right under this Agreement constitute a continuing waiver of the same or a waiver of any other right under this Agreement.

ARTICLE 11 – SEVERABILITY

11.1 Severability. If any term of this Agreement is invalid, illegal, or unenforceable in any jurisdiction, such invalidity, illegality, or unenforceability shall not affect any other term of this Agreement or invalidate or render unenforceable the term in any other jurisdiction. Upon a determination that any term is invalid, illegal, or unenforceable, the Parties shall negotiate in good faith to modify this Agreement to reflect the original intent of the Parties as closely as possible in a mutually acceptable manner in order that the activities contemplated by this Agreement be performed to the greatest extent possible as originally contemplated.

ARTICLE 12 – GOVERNING LAW AND JURISDICTION

12.1 Governing Law. This Agreement shall be governed by, and constructed and interpreted in accordance with the laws of the United States.

12.2 Submission to Jurisdiction. The Parties irrevocably submit to the exclusive jurisdiction of the federal courts located in Washington, District of Columbia for any legal suit, action, or proceeding arising out of, based upon, or relating to this Agreement. Notwithstanding the foregoing, the Department does not waive its sovereign immunity.

ARTICLE 13 – NOTICES AND CONTACT INFORMATION

13.1 Points of Contact. The Parties will provide each other with contact details for the representative with responsibility for addressing day-to-day inquiries no later than five days after execution of the Agreement. Each Party may unilaterally change its point of contact with one day written notice to the other Party.

13.2 Notices. Any notices sent under this Agreement must be in writing and sent via email to the individual identified pursuant to Section 13.1, Points of Contact.

ARTICLE 14 – RESPONSIBILITY FOR STAFF

14.1 Independent Parties. The relationship of the Parties to this Agreement is that of independent parties and not as agents of each other or as joint venturers or partners. Each Party shall maintain sole and exclusive control over its personnel and operations.

14.2 Responsibility for Staff. The Collaborator will provide adequate supervision of and care for its staff, authorized agents, and representatives whenever they are carrying out activities related to this agreement.

ARTICLE 15 – ASSIGNMENT AND NOVATION

15.1 Assignment and Novation. The Collaborator shall not assign any of its rights, delegate any of its obligations, or overwise execute a novation with respect to this Agreement without the prior written consent of CAISI.

ARTICLE 16 – COSTS

16.1 Costs. The Collaborator understands and acknowledges that this is a no-cost Agreement and it will not receive any financial compensation for performance under the Agreement. Each Party will be responsible for its own costs incurred in connection with this Agreement. The Collaborator will bear any costs associated with hosting and running its AI Systems for the duration of the Agreement.

16.2 Antideficiency Act. The Parties performance under this Agreement is subject to the availability of funds. Neither CAISI nor its employees or officers may create or authorize an obligation in excess of the funds available or in advance or in lieu of appropriations unless otherwise authorized by law.

ARTICLE 17 – AMENDMENT OF THIS AGREEMENT

17.1 Amendment of the Agreement. If either Party desires an amendment of this Agreement, the Parties shall, upon reasonable notice of the proposed amendment by the Party desiring the change, confer in good faith to determine the desirability of such amendment. Such amendment shall not be effective until a written amendment is signed by all Parties hereto by their representatives duly authorized to execute such amendment.

ARTICLE 18 – ENTIRETY OF AGREEMENT

18.1 Entirety of Agreement. This Agreement constitutes the entire agreement between the Parties and supersedes all prior and contemporaneous agreements, understandings, negotiations, and discussions amongst the Parties, whether oral or written, with respect to the subject matter hereof. In the event of any conflict between the terms of this Agreement and an Evaluation Plan, this Agreement will control. Notwithstanding the preceding sentence, with respect to any given evaluation, if a provision in the respective Evaluation Plan expressly states that such provision will control over a specific provision of this Agreement, then such provision in the Evaluation Plan will control.

ARTICLE 19 – SIGNATORIES

19.1 Signatories. By signing this Agreement, the signatories represent and warrant that they have the authority to act on behalf of their respective entities in entering into this Agreement:

For the Collaborator:
Open AI OpCo LLC

[redacted]

Date: 7/31/2025

For the Government:
NIST, CAISI

[redacted]

Date: 7/31/2025