PALANTIR AIP ADDENDUM

The customer ("**Customer**") agreeing to the terms of this Palantir AIP Addendum (the "**Addendum**") and the Palantir Technologies entity that is the signatory to the Agreement ("**Palantir**"; each a "**Party**" and collectively the "**Parties**"), have entered into an agreement governing Customer's use of the Service (the "**Agreement**"). This Addendum is incorporated into the Agreement. Any capitalized terms used but not defined in this Addendum shall have the meaning provided in the Agreement. If there is a conflict in meaning between the Agreement and this Addendum, this Addendum shall prevail and control. This Addendum will remain in effect until any expiration or termination of the Agreement pursuant thereto.

1. Certain Definitions.

1.1 "**AIP**" or "Palantir Artificial **Intelligence** Platform" means Palantir's proprietary software-as-a-service offering(s) leveraging artificial intelligence (including but not limited to language models and other modeling services) ("**Models**"), including the Third Party Model Services set forth in Exhibit A attached hereto, and as set forth in the Documentation and/or the Agreement.

1.2 "**Documentation**" means any technical documentation for the Service made available in connection with the Service.

1.3 "**Service**" means Palantir's proprietary software-as-a-service offering(s) set forth in the Agreement.

2. Customer use of AIP.

2.1 AIP Use. AIP employs artificial intelligence and machine learning techniques, including use of Models. Customer acknowledges that because of the statistical methods underlying the foregoing techniques, output of AIP may be incorrect, incomplete, or biased. Accordingly, Customer agrees that its use of AIP is at its own risk and Customer shall be solely responsible for any decisions, actions, or omissions it takes on the basis of any AIP output. Customer further agrees to evaluate (including through review by a natural person) any AIP output prior to taking any decisions, actions, or omissions on its basis. Access to AIP forms part of the Service and is subject to relevant terms and conditions of the Agreement applicable to the Service, and Customer's right and/or license (as applicable) to use AIP is determined by its right and/or license to use and/or access the Service as provided in the Agreement. Customer may only use data in connection with AIP for which Customer and/or its user has received all consents, authorizations, approvals, and/or agreements necessary to permit such use and/or processing under applicable law, regulation, or agreement(s). Customer shall not use AIP to attempt to obtain any information that may violate third party rights or applicable laws or regulations, including classified information. Customer's access and use of AIP shall comply with all applicable laws and regulations. Customer shall be fully responsible and liable for its users' use of and access to AIP. Customer's use of AIP leveraging a Third Party Model Service (as defined in Exhibit A, attached hereto) shall comply with the Additional Terms set forth in Exhibit A, applicable to such Third Party Model Service.

2.2 Restrictions. Customer will not (and will not allow any third party to): (a) decompile, disassemble, scan, reverse engineer, or attempt to discover any source code, algorithms, weights of the underlying models, or underlying ideas of AIP; or (b) represent that any output from AIP was generated by a natural person when it was not.

2.3 Feedback. Customer may share and provide feedback, suggestions, comments, ideas, data, including prompts, reports, or other information related to Customer's use of AIP in the form of feedback report(s) to Palantir ("**AIP Feedback**") and Palantir can directly request AIP Feedback from Users of AIP. Customer acknowledges that any AIP Feedback shared to Palantir is done so voluntarily. Customer grants to Palantir a worldwide, perpetual, irrevocable, and royalty-free right and license to use, distribute, disclose, and make and incorporate into AIP any AIP Feedback to provide, secure, and improve AIP.

3. Fees. Customer's use of AIP shall accrue fees on the same terms as Customer's use of the Service pursuant to the Agreement. Notwithstanding the foregoing, if the Agreement does not include separate fees for use of AIP, Palantir reserves the right to begin charging for use of AIP upon providing Customer thirty (30) days' notice to opt in or otherwise lose access to AIP.

4. Suspension of Services. If Palantir reasonably determines that: (a) Customer's access or use of any Model violates applicable law, regulation, or any material term of the Agreement, including this Addendum; (b) Palantir providing any part of AIP would violate applicable (in force or forthcoming) law, regulation, or agreements; or (c) Palantir's provision of any part of AIP poses undue security risk to Palantir or its customers, Palantir reserves the right to disable, suspend, or terminate Customer's access to all or any part of AIP. Palantir will notify Customer prior to exercising the foregoing right concurrent or prior to such exercise.

5. Disclaimer. NOTWITHSTANDING ANYTHING TO THE CONTRARY IN THE AGREEMENT, AIP IS PROVIDED "AS-IS" WITHOUT ANY OTHER WARRANTIES OF ANY KIND AND PALANTIR HEREBY DISCLAIMS ALL WARRANTIES OF ANY KIND, WHETHER EXPRESS OR IMPLIED, ORAL OR WRITTEN, INCLUDING BUT NOT LIMITED TO ANY WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY, TITLE, SATISFACTORY QUALITY, OR FITNESS FOR A PARTICULAR PURPOSE. WITHOUT LIMITING THE FOREGOING LIMITATION, PALANTIR DOES NOT WARRANT THAT AIP WILL MEET CUSTOMER REQUIREMENTS OR GUARANTEE ANY QUALITY, RESULTS, OUTCOMES, OR CONCLUSIONS OR THAT OPERATION OF AIP WILL BE UNINTERRUPTED OR ERROR FREE. PALANTIR IS NOT RESPONSIBLE FOR ANY DECISIONS, ACTIONS, OR OMISSIONS CUSTOMER TAKES BASED UPON OR INFORMED BY OUTPUT FROM AIP. PALANTIR IS NOT RESPONSIBLE OR LIABLE FOR ANY THIRD PARTY SERVICES PROVIDED IN THE COURSE OF DELIVERING AIP, INCLUDING (IF APPLICABLE) ANY THIRD PARTY MODEL SERVICE (INCLUDING WITHOUT LIMITATION, UPTIME GUARANTEES, OUTAGES, FAILURES, OR ANY OTHER GUARANTEES IN ANY SERVICE LEVEL AGREEMENT BETWEEN THE PARTIES), CUSTOMER'S INPUT TO AIP, OR OUTPUT FROM AIP (INCLUDING BUT NOT LIMITED TO COMPLETENESS, OR ACCURACY OF OUTPUT FROM AIP, OR WHETHER THE OUTPUT FROM AIP INFRINGES OR VIOLATES ANY THIRD PARTY'S RIGHTS, INCLUDING INTELLECTUAL PROPERTY AND CONTRACTUAL RIGHTS).

6. Miscellaneous. This Addendum and any dispute or claim (including any non-contractual disputes or claims) arising out of or in connection with it, or its subject matter or formation, shall be governed by and construed in accordance with the laws that govern the Agreement and the dispute resolution provisions therein. If any provision of this Agreement shall be adjudged by any court of competent jurisdiction to be unenforceable or invalid, that provision shall be limited or eliminated to the minimum extent necessary so that this Agreement shall otherwise remain in full force and effect and be enforceable. No waiver of any breach shall be deemed a waiver of any subsequent breach.

EXHIBIT A

Customer's use of AIP may leverage the following Models hosted in a third party environment (each a "Third Party Model Service").

Third Party Model Service	**Additional Terms**
Anthropic Models hosted in Palantir's Anthropic Environment ("**Anthropic Model Service**")	Customer's use of AIP leveraging Anthropic Model Service (a) shall comply with the Anthropic Usage Policy (available at https://www.anthropic.com/legal/aup); and (b) shall not facilitate or engage in the following: (i) design, market, help distribute or utilize weapons, explosives, dangerous materials or other systems designed to cause harm to or loss of human life; (ii) covertly tracking, targeting, or surveilling individuals, i.e., searching for or gathering information on an individual or group in order to track, target or report on their identity, including using the product for facial recognition, covert tracking, battlefield management applications or predictive policing; (iii) automated determination of financing eligibility of individuals, i.e., making automated decisions about the eligibility of individuals for financial products and creditworthiness; (iv) automated determination of employment and housing decisions, i.e., making automated decisions about the employability of individuals or other employment determinations or

	decisions regarding eligibility for housing, including leases and home loans; (v) any law enforcement application, except for the following permitted applications by U.S. law enforcement organizations: back office uses including call center support, document summarization, and accounting; (vi) analysis of data for the location of missing persons and other applications, provided that such applications do not otherwise violate or impair the liberty, civil liberties, or human rights of natural persons; (vii) shall not be used for the purpose of building a product or service competitive to the Anthropic Model Service (including to train competing Models or to resell the Models); and (viii) shall not involve reverse engineering or duplicating the Anthropic Model Service. Customer acknowledges that Anthropic, PBC may collect and temporarily retain pseudonymized security classifier metadata related to Customer's use of AIP leveraging the Anthropic Model Service (which metadata, for the avoidance of doubt, shall not include the contents of Customer's prompts provided to or output received from the Anthropic Model Service).
Models hosted in Palantir's Microsoft Azure Environment ("**Azure Model Service**")	(a) Customer's use of AIP leveraging OpenAI Models via the Azure Model Service shall comply with the Azure OpenAI Code of Conduct (https://learn.microsoft.com/en-us/legal/cognitive-services/openai/code-of-conduct?context=%2Fazure%2Fcognitive-services%2Fopenai%2Fcontext%2Fcontext); (b) Customer shall only use AIP leveraging the Azure Model Service to (i) submit content to be summarized for pre-defined topics built into AIP and cannot use AIP as an open-ended summarizer (examples of such permitted use include but are not limited to summarization of call center transcripts, technical reports, and product reviews); (ii) analyze inputs using classification, sentiment analysis of text, or entity extraction (examples of such permitted use include but are not limited to analyzing product feedback sentiment, analyzing support calls and transcripts, and refining text-based search with embeddings; (iii) search trusted source documents such as internal Customer documentation; (iv) ask questions and receive answers from trusted source documents such as internal Customer documentation; or (v) code generation or transformation scenarios (examples of such permitted use include but are not limited to converting one programming language to another, generating docstrings for functions, or converting natural language to SQL); (c) Customer shall not use AIP leveraging OpenAI Models via the Azure Model Service (i) to generate, distribute, or modify any output from the Azure Model Service that the Customer knew or should have known was infringing or likely to infringe a third party's intellectual property or other proprietary rights (including if such infringement is caused by Customer's combination of such output with third party products or services); (ii) use output from the OpenAI Models via the Azure Model Service to develop any artificial intelligence models that compete with OpenAI's products and services; or (iii) while disabling, ignoring, or otherwise circumventing, without authorization, any relevant citation, filtering, or safety features or restrictions provided by Azure or Palantir applicable to the Azure Model Service; (d) Customer's use of AIP leveraging Anthropic Models via the Azure Model Service (i) shall comply with the Anthropic Usage Policy (available at https://www.anthropic.com/legal/aup); (ii) shall not be used for the purpose of building a product or service competitive to the Anthropic Models (including to train competing Models or to resell the Models); and (iii) shall not involve reverse engineering or duplicating the Anthropic Models; and (e) Customer's use of AIP leveraging NVIDIA Models via the Azure Model Service shall comply with the applicable NVIDIA license terms (available at https://ai.azure.com/catalog/models/ for the applicable NVIDIA models, such as those listed at https://ai.azure.com/catalog/models/NVIDIA-Nemotron-3-Ultra-NIM-microservice for the NVIDIA-Nemotron-3-Ultra-NIM-microservice model). Customer acknowledges that Microsoft Corporation may collect and temporarily retain pseudonymized security classifier metadata related to Customer's use of AIP leveraging Anthropic Models via the Azure

	Model Service (which metadata, for the avoidance of doubt, shall not include the contents of Customer's prompts provided to or output received from the Azure Model Service).
Models hosted in Palantir's Amazon Web Services Environment ("**AWS Model Service**")	Customer's use of AIP leveraging Anthropic Models through the AWS Model Service (a) shall comply with the Anthropic Code of Conduct (https://console.anthropic.com/legal/aup), (b) shall comply with the Anthropic Bedrock AI Services Agreement (available at https://s3.amazonaws.com/EULA/Anthropic-EULA-1023.pdf), and (c) shall not facilitate or engage in the following: (i) design, market, help distribute or utilize weapons, explosives, dangerous materials or other systems designed to cause harm to or loss of human life; (ii) covertly tracking, targeting, or surveilling individuals, i.e., searching for or gathering information on an individual or group in order to track, target or report on their identity, including using the product for facial recognition, covert tracking, battlefield management applications or predictive policing; (iii) automated determination of financing eligibility of individuals, i.e., making automated decisions about the eligibility of individuals for financial products and creditworthiness; (iv) automated determination of employment and housing decisions, i.e., making automated decisions about the employability of individuals or other employment determinations or decisions regarding eligibility for housing, including leases and home loans; (v) any law enforcement application, except for the following permitted applications by U.S. law enforcement organizations: back office uses including call center support, document summarization, and accounting; (vi) analysis of data for the location of missing persons and other applications, provided that such applications do not otherwise violate or impair the liberty, civil liberties, or human rights of natural persons; (vii) use for the purpose of building a product or service competitive to the Anthropic Models (including to train competing Models or to resell the Models); or (viii) reverse engineering or duplicating the Anthropic Models. Customer's use of NVIDIA Models through the AWS Model Service shall comply with the NVIDIA Open Model License Agreement (https://www.nvidia.com/en-us/agreements/enterprise-software/nvidia-open-model-license/), which includes NVIDIA's Trustworthy AI terms incorporated by reference (https://www.nvidia.com/en-us/agreements/trustworthy-ai/terms/). Customer hereby agrees that its use of AIP leveraging Models through the AWS Model Service shall comply with any acceptable use policies or codes of conduct applicable to such Models, as made available to Customer through AIP or the Documentation. Customer's use of AIP leveraging OpenAI Models via the AWS Model Service (1) shall comply with the OpenAI Services Agreement – Amazon Bedrock (available at https://aws.amazon.com/legal/bedrock/third-party-models/#amsc13--xttmgl); and (2) shall not involve use of output from the OpenAI Models via the AWS Model Service to develop any artificial intelligence models that compete with OpenAI's products and services. Customer acknowledges that Amazon Web Services, Inc. may collect and temporarily retain pseudonymized security classifier metadata related to Customer's use of AIP leveraging the AWS Model Service (which metadata, for the avoidance of doubt, shall not include the contents of Customer's prompts provided to or output received from the AWS Model Service).
Google Provided Models hosted in Palantir's Google Cloud Services	Customer's use of AIP leveraging the Google Model Service (a) shall comply with the Google Generative AI Prohibited Use Policy (https://policies.google.com/terms/generative-ai/use-policy); (b) shall comply with the Google Cloud Platform Acceptable Use Policy (https://cloud.google.com/terms/aup); (c) with respect to Anthropic Models used or accessed via the Google Model Service, Customer's use of AIP leveraging such models (i) shall comply with the

Environment (“**Google Model Service**”)	Anthropic on Vertex - Commercial Terms of Service (available at https://www-cdn.anthropic.com/471bd07290603ee509a5ea0d5ccf131ea5897232/anthropic-vertex-commercial-terms-march-2024.pdf); (ii) shall not be used for the purpose of building a product or service competitive to the Anthropic Models (including to train competing Models or to resell the Models); and (iii) shall not involve reverse engineering or duplicating the Anthropic Models; (d) shall not reasonably be expected to lead to death, personal injury, or environmental damage, including operation of nuclear facilities, air traffic control, life support systems, or weaponry; and (e) shall, with respect to Customer’s use of applicable models that meet the definition of “Pre-GA Offerings” in the Google Model Service Pre-GA Offerings Terms (defined below) made available via the Google Model Service, comply with Google’s “Pre-GA Offerings Terms” subsection in the “General Service Terms” section of the Google Cloud Platform Service Specific Terms, available at https://cloud.google.com/terms/service-terms (the “Google Model Service Pre-GA Offerings Terms”). Customer acknowledges that Google LLC may collect and temporarily retain pseudonymized security classifier metadata related to Customer’s use of AIP leveraging the Google Model Service (which metadata, for the avoidance of doubt, shall not include the contents of Customer’s prompts provided to or output received from the Google Model Service).
OpenAI Models hosted by OpenAI (“**OpenAI Model Service**”)	Customer’s use of AIP leveraging the OpenAI Model Service (a) shall comply with the OpenAI Usage Policies (https://openai.com/policies/usage-policies); (b) if and only as applicable, shall comply with the applicable OpenAI Service Terms (https://openai.com/policies/service-terms); (c) Customer shall only use AIP leveraging the OpenAI Model Service to (i) submit content to be summarized for pre-defined topics built into AIP and cannot use AIP as an open-ended summarizer (examples of such permitted use include but are not limited to summarization of call center transcripts, technical reports, and product reviews); (ii) analyze inputs using classification, sentiment analysis of text, or entity extraction (examples of such permitted use include but are not limited to analyzing product feedback sentiment, analyzing support calls and transcripts, and refining text-based search with embeddings; (iii) search trusted source documents such as internal Customer documentation; (iv) ask questions and receive answers from trusted source documents such as internal Customer documentation; (v) code generation or transformation scenarios (examples of such permitted use include but are not limited to converting one programming language to another, generating docstrings for functions, or converting natural language to SQL); or (vi) fine-tune Models as provided as part of AIP and the OpenAI Model Service; and (d) Customer shall not use AIP leveraging the OpenAI Model Service (i) to generate, distribute, or modify any output from the OpenAI Model Service that the Customer knew or should have known was infringing or likely to infringe a third party’s intellectual property or other proprietary rights (including if such infringement is caused by Customer’s combination of such output with third party products or services); (ii) use output from the OpenAI Models via the Azure Model Service to develop any artificial intelligence models that compete with OpenAI’s products and services; or (iii) while disabling, ignoring, or otherwise circumventing, without authorization, any relevant citation, filtering, or safety features or restrictions provided by OpenAI applicable to the OpenAI Model Service. Customer acknowledges that OpenAI, LLC may collect and temporarily retain pseudonymized security classifier metadata related to Customer’s use of AIP leveraging the OpenAI Model Service (which metadata, for the avoidance of doubt, shall not include the contents of Customer’s prompts provided to or output received from the OpenAI Model Service).
Meta’s Llama 4 Model, Meta's Llama 3.3	Customer’s use of AIP leveraging the Llama Model Service shall comply with (a) the Llama 4 Community License Agreement (https://www.llama.com/llama4/license/), which includes the Llama 4

Model, Meta's Llama 3.2 Model, and Meta's Llama 3.1 Model (and Meta's Llama Models now existing or released in the future) hosted in Palantir's model hub environment(s) ("**Llama Model Service**")	Acceptable Use Policy incorporated by reference (https://www.llama.com/llama4/use-policy/); (b) the Llama 3.3 Community License Agreement (https://www.llama.com/llama3_3/license/), which includes the Llama 3.3 Acceptable Use Policy incorporated by reference (https://www.llama.com/llama3_3/use-policy/); (c) the Llama 3.2 Community License Agreement (https://github.com/meta-llama/llama-models/blob/main/models/llama3_2/LICENSE), which includes the Llama 3.2 Acceptable Use Policy (https://www.llama.com/llama3_2/use-policy/); (d) the Llama 3.1 Community License Agreement (https://github.com/meta-llama/llama-models/blob/main/models/llama3_1/LICENSE), which includes the Llama 3.1 Acceptable Use Policy (https://llama.meta.com/llama3_1/use-policy/); and (e) any license terms and/or acceptable use policies applicable to any other Llama Model(s) (now existing or released in the future) made available as part of the Llama Model Service.
X.AI LLC's Grok Models hosted by X.AI LLC ("**xAI Model Service**")	Customer's use of AIP leveraging the xAI Model Service shall (a) comply with applicable law (this includes not taking unlawful action on behalf of others); (b) not promote or engage in (i) the sexualization or exploitation of children, (ii) violating copyright, trademark, or other intellectual property law, (iii) compromising others' privacy (iv) violating a person's right to publicity, (v) operating in a regulated industry without complying with those regulations, or (vi) defrauding, defaming, scamming, or spamming; (c) not harm people or property (this includes but is not limited to not perpetrating hacking, doxing, phishing or other malicious attacks, and not developing any kind of weapon); (d) not mislead or circumvent technical safeguards or other technical guardrails; (e) not represent that output of the xAI Model Service is human generated; and (f) not use output of the xAI Model Service to train generative artificial intelligence models similar to or competitive with the xAI Model Service.
NVIDIA's llama-3.3-Nemotron-Super-49B-v1 model and NVIDIA's llama-3.2-nv-embedqa-1b-v2 model hosted in Palantir's model hub environment(s) ("**NVIDIA Model Service**")	Customer's use of AIP leveraging the NVIDIA Model Service shall comply with (a) the NVIDIA Open Model License Agreement (https://www.nvidia.com/en-us/agreements/enterprise-software/nvidia-open-model-license/), which includes NVIDIA's Trustworthy AI terms incorporated by reference (https://www.nvidia.com/en-us/agreements/trustworthy-ai/terms/); (b) the Llama 3.3 Community License Agreement (https://www.llama.com/llama3_3/license/), which includes the Llama 3.3 Acceptable Use Policy incorporated by reference (https://www.llama.com/llama3_3/use-policy/); (c) the Llama 3.2 Community License Agreement (https://github.com/meta-llama/llama-models/blob/main/models/llama3_2/LICENSE), which includes the Llama 3.2 Acceptable Use Policy (https://www.llama.com/llama3_2/use-policy/); and (d) any license terms and/or acceptable use policies applicable to any other NVIDIA Model(s) (now existing or released in the future) made available as part of the NVIDIA Model Service.