Commercial Terms of Service

These Commercial Terms of Service ("**Terms**") are an agreement between Anthropic, PBC and the Ordering Activity under GSA Schedule contracts ("**Customer**") identified in the order between Customer and Carahsoft Technology Corporation ("**Reseller**" and such order, the "**Order**"). They govern Customer's use of the Anthropic offerings referenced in the Order and subsequent orders, as well as all related Anthropic tools, documentation and services (the "**Services**"). These Terms are effective on the start date set forth in Customer's first Order ("**Effective Date**").

A. **Services.**

1. **Overview.** Subject to these Terms, Anthropic gives Customer permission to use the Services, including to power products and services Customer makes available to its own customers and end users ("**Users**").

2. **Third Party Features.** Customer may elect (in its sole discretion) to use features, services or other content may be made available by third parties to Customer through the Services ("**Third Party Features**"). Customer acknowledges and agrees that Third Party Features are not Services and, accordingly, Anthropic is not responsible for them.

3. **Feedback.** If Customer provides (in its sole discretion) Anthropic with feedback regarding the Services, Anthropic may use that feedback at its own risk and without obligation to Customer. Anthropic acknowledges that these Terms and any Feedback provided under these Terms in advertising is limited by GSAR 552.203-7.

B. **Customer Content.** As between the parties and to the extent permitted by applicable law, Anthropic agrees that Customer (a) retains all rights to its Inputs, and (b) owns its Outputs. Anthropic disclaims any rights it receives to the Customer Content under these Terms. Subject to Customer's compliance with these Terms, Anthropic hereby assigns to Customer its right, title and interest (if any) in and to Outputs. Anthropic may not train models on Customer Content from Services. "**Inputs**" means submissions to the Services by Customer or its Users and "**Outputs**" means responses generated by the Services to Inputs (Inputs and Outputs together are "**Customer Content**").

C. **Data Privacy.** Data submitted through the Services will be processed in accordance with the Anthropic Data Processing Addendum ("**DPA**"), which will be attached to the Order with Reseller and and in effect as of the Order start date of the applicable Order with Reseller. Notwithstanding the terms in the DPA, If Customer is an instrumentality of the United States, Customer will only be required to comply with the Federal laws of the United States for any Customer obligations or responsibilities set forth in the DPA.

D. **Trust and Safety; Restrictions.**

1. **Compliance.** Each party will comply with all laws applicable to the provision (for Anthropic) and use (for Customer) of the Services, including any applicable data privacy laws.

2. **Policies and Service Terms.** Customer and its Users may only use the Services in compliance with these Terms, including (a) the Usage Policy ("**Usage Policy**", which was previously referred to as the Acceptable Use Policy), (b) our policy on the countries and regions Anthropic currently supports ("**Supported Regions Policy**") and (c) our Service Specific Terms, which will be attached to the applicable Order with Reseller and in effect as of the Order start date of the applicable Order with Reseller. Customer must cooperate with reasonable requests for information from Anthropic to support compliance with its Usage Policy, including to verify Customer's identity and use of the Services.

3. **Limitations of Outputs; Notice to Users.** It is Customer's responsibility to evaluate whether Outputs are appropriate for Customer's use case, including where human review is appropriate, before using or sharing Outputs. Customer acknowledges, and must notify its Users, that factual assertions in Outputs should not be relied upon without independently checking their accuracy, as they may be false, incomplete, misleading or not reflective of recent events or information. Customer further acknowledges that Outputs may contain content inconsistent with Anthropic's views.

4. **Use Restrictions.** Customer may not and must not attempt to (a) access the Services to build a competing product or service, including to train competing AI models or resell the Services except as expressly approved by Anthropic; (b) reverse engineer or duplicate the Services; or (c) support any third party's attempt at any of the conduct restricted in this sentence.

5. **Service Account.** Customer is responsible for all activity under its account. Customer will promptly notify Anthropic if Customer believes the account it uses to access the Services has been compromised, or is subject to a denial of service or similar malicious attack that may negatively impact the Services.

E. **Confidentiality.**

1. **Confidential Information.** The parties may share information that is identified as confidential, proprietary, or similar, or that a party would reasonably understand to be confidential or proprietary (**"Confidential Information"**). Customer Content is Customer's Confidential Information.

2. **Obligations of Parties.** The receiving party (**"Recipient"**) may only use Confidential Information of the disclosing party (**"Discloser"**) to exercise its rights and perform its obligations under these Terms. Recipient may only share Discloser's Confidential Information to Recipient's employees, agents, and advisors that have a need to know such Confidential Information and who are bound to obligations of confidentiality at least as protective as those provided in these Terms (**"Representatives"**). Recipient will protect Discloser's Confidential Information from

unauthorized use, access, or disclosure in the same manner as Recipient protects its own Confidential Information, and with no less than reasonable care. Recipient is responsible for all acts and omissions of its Representatives.

3. **Exclusions.** Confidential Information excludes information that: (a) becomes publicly available through no fault of Recipient; (b) is obtained by Recipient from a third party without a breach of the third party's obligations of confidentiality; or (c) is independently developed by Recipient without use of Confidential Information. Recipient may disclose Discloser's Confidential Information to the extent it is required by law, or court or administrative order, and will, except where expressly prohibited, notify Discloser of the required disclosure promptly and fully cooperate with Discloser's efforts to prevent or narrow the scope of disclosure. Anthropic recognizes that Federal agencies are subject to the Freedom of Information Act, 5 U.S.C. 522("FOIA"), which may require that certain information be released, despite being characterized as "confidential" by Anthropic. Customer will provide Anthropic with prior written notice (to the extent legally permitted) before disclosing any Anthropic Confidential Information, allowing Anthropic to seek available exemptions under FOIA.

4. **Destruction Request.** Recipient will destroy Discloser's Confidential Information promptly upon request, except where retained to comply with law or copies in Recipient's automated back-up systems, which will remain subject to these obligations of confidentiality while maintained.

F. **Intellectual Property.** Except as expressly stated in these Terms, these Terms do not grant either party any rights to the other's content or intellectual property, by implication or otherwise.

G. **Publicity.** Anthropic may use Customer's name and logo to publicly identify Customer as a customer of the Services; provided that Customer may opt-out via this request form. Customer will consider in good faith any request by Anthropic to (1) provide a quote from a Customer executive regarding Customer's motivation for using the Services that Anthropic may use publicly and (2) participate in a public co-marketing activity.

H. **Fees.**

1. **Payment of Fees.** Customer is responsible for fees incurred by its account, at the rates specified in the applicable Order with Reseller.

2. **Reserved.**

3. **Reserved.**

I. **Termination and Suspension.**

1. **Term.** These Terms start on the Effective Date and continue until terminated (the "**Term**").

2. **Termination.**

a. As permitted by applicable law, including the Contracts Disputes Act (41 U.S.C. SS7101-7109) where Customer is an instrumentality of the U.S. Federal Government subject to the Contracts Disputes Act, Customer's use of the Services may be immediately terminated and/or suspended, at Anthropic's option, upon notice due to (a) a breach of these Terms, or Orders by Customer or any User; or (b) a breach by Reseller of Reseller's payment obligations to Anthropic with respect to the Services it is reselling to Customer.

b. The parties acknowledge that Anthropic is under no obligation to (i) continue providing the Services directly to Customer; or (ii) assume a direct contractual relationship with Customer following the termination or expiration of Reseller's agreement with Anthropic.

c. **Reserved**.

3. **Suspension.**

a. Notwithstanding Section I(2)(a), Anthropic may temporarily suspend Customer's access to any portion or all of the Services if Anthropic reasonably believes or determines that there is a risk to or attack on any of the Services (a "**Service Suspension**").

b. Anthropic will use reasonable efforts to provide written notice of any Service Suspension to Customer, and resume providing access to the Services, as soon as reasonably possible after the event giving rise to the Service Suspension is cured, where curable. Anthropic will have no liability for any damage, liabilities, losses (including any loss of data or profits), or any other consequences that Customer may incur because of a Service Suspension.

4. **Effect of Termination.** Upon termination, Customer may no longer access the Services. The following provisions will survive termination or expiration of these Terms: (a) Sections E (Confidentiality), G (Publicity), H (Fees), I (Termination and Suspension), J (Disputes), K (Indemnification), L.2 (Disclaimer of Warranties), L.3 (Limits on Liability), and M (Miscellaneous); (b) any provision or condition that must survive to fulfill its essential purpose.

J. **Disputes.**

1. **Disputes.** When Customer is an instrumentality of the U.S. Federal Government subject to the Contract Disputes Act , any claims against Customer for alleged breach of these Terms must be brought under the Contract Disputes Act. Anthropic will continue performing under these Terms, pending final resolution of any claim, appeal, or action arising under these Terms, and comply with any decision of the Contracting Officer.

2. **Reserved.**

3. **Reserved.**

K. **Indemnification.**

1. **Claims Against Customer.** Anthropic will have the right to intervene to defend Customer and its personnel, successors, and assigns from and against any Customer Claim (as defined below) and indemnify them for any judgment that a court of competent jurisdiction grants a third party on such Customer Claim or that an arbitrator awards a third party under any Anthropic-approved settlement of such Customer Claim. "**Customer Claim**" means a third-party claim, suit, or proceeding alleging that Customer's paid use of the Services (which includes data Anthropic has used to train a model that is part of the Services) in accordance with these Terms or Outputs generated through such authorized use violates any third-party intellectual property right. Nothing contained herein will be construed in derogation of the U.S. Department of Justice's right to defend any claim or action brought against the U.S., pursuant to its jurisdictional statute 28 U.S.C. §516.

2. **Claims Against Anthropic.** This Section K(2) will only be enforceable as permitted by applicable law and will not apply if Customer is an instrumentality of the United States Federal Government that is subject to the Anti-Deficiency Act (31 U.S.C. 1341). Customer will defend Anthropic and its personnel, successors, and assigns from and against any Anthropic Claim (as defined below) and indemnify them for any judgment that a court of competent jurisdiction grants a third party on such Anthropic Claim or that an arbitrator awards a third party under any Customer-approved settlement of such Anthropic Claim. "**Anthropic Claim**" means any third-party claim, suit, or proceeding related to Customer's or its Users' (a) Inputs or other data provided by Customer, or (b) use of the Services in violation of the Usage Policy, the Service Specific Terms, or Section D.4 (Use Restrictions). Anthropic Claims and Customer Claims are each a "**Claim**", as applicable.

3. **Exclusions.** Neither party's defense or indemnification obligations will apply to the extent the underlying allegation arises from the indemnified party's fraud, willful misconduct, violations of law, or breach of the Agreement. Additionally, Anthropic's defense and indemnification obligations will not apply to the extent the Customer Claim arises from: (a) modifications made by Customer to the Services or Outputs; (b) the combination of the Services or Outputs with technology or content not provided by Anthropic; (c) Inputs or other data provided by Customer; (d) use of the Services or Outputs in a manner that Customer knows or reasonably should know violates or infringes the rights of others; (e) the practice of a patented invention contained in an Output; or (f) an alleged violation of trademark based on use of an Output in trade or commerce.

4. **Process.** The indemnified party must promptly notify the indemnifying party of the relevant Claim, and will reasonably cooperate in the defense. The indemnifying party will retain the right to control the defense of any such Claim, including the selection of counsel, the strategy and course of any litigation or appeals, and any negotiations or settlement or compromise, except that the indemnified party will have the right, not to be exercised unreasonably, to reject any settlement or compromise that requires that it admit wrongdoing or liability or subjects it to an ongoing affirmative obligation. The indemnifying party's obligations will be excused if either of the following materially prejudices the defense: (a) failure of the indemnified party to provide prompt notice of the Claim; or (b) failure to reasonably cooperate in the defense.

5. **Sole Remedy.** To the extent covered under this Section K (Indemnification), indemnification is each party's sole and exclusive remedy under these Terms for any third-party claims.

L. **Warranties and Limits on Liability.**

1. **Warranties.** Each party represents and warrants that (a) it is authorized to enter into these Terms; and (b) entering into and performing these Terms will not violate any of its corporate rules, if applicable. Customer further represents and warrants that it has all rights and permissions required to submit Inputs to the Services. Anthropic warrants that during the applicable Order term that (i) the Services have been developed and are intended to operate in accordance with the applicable model card(s) found at https://docs.anthropic.com/en/docs/about-claude/models and (ii) the overall security of the Services will not materially decrease.

2. **Disclaimer of Warranties.** EXCEPT TO THE EXTENT EXPRESSLY PROVIDED FOR IN THESE TERMS, TO THE MAXIMUM EXTENT PERMITTED UNDER LAW (A) THE SERVICES AND OUTPUTS ARE PROVIDED "AS IS" AND "AS AVAILABLE" WITHOUT WARRANTY OF ANY KIND; AND (B) ANTHROPIC MAKES NO WARRANTIES, EXPRESS OR IMPLIED, RELATING TO THIRD-PARTY PRODUCTS OR SERVICES, INCLUDING THIRD-PARTY INTERFACES. ANTHROPIC EXPRESSLY DISCLAIMS ALL IMPLIED WARRANTIES, INCLUDING WARRANTIES OF MERCHANTABILITY, NON-INFRINGEMENT, AND FITNESS FOR A PARTICULAR PURPOSE, AS WELL AS ANY IMPLIED WARRANTY ARISING FROM STATUTE, COURSE OF DEALING OR PERFORMANCE, OR TRADE USE. ANTHROPIC DOES NOT WARRANT, AND DISCLAIMS THAT, THE SERVICES OR OUTPUTS ARE ACCURATE, COMPLETE OR ERROR-FREE OR THAT THEIR USE WILL BE UNINTERRUPTED. REFERENCES TO A THIRD PARTY IN THE OUTPUTS MAY NOT MEAN THEY ENDORSE OR ARE OTHERWISE WORKING WITH ANTHROPIC.

3. **Limits on Liability.**

a. Except as stated in Section L.3.b, the liability of each party, and its affiliates and licensors, for any damages arising out of or related to these Terms (i) excludes damages that are consequential, incidental, special, indirect, or exemplary damages, including lost profits, business, contracts, revenue, goodwill, production, anticipated savings, or data, and costs of procurement of substitute goods or services and (ii) is limited to Fees paid by Reseller to Anthropic for Customer's use the Services in the previous 12 months.

b. The limitations of liability in this Section L.3 (Limits on Liability) do not apply to either party's obligations under Section K (Indemnification).

c. THE LIMITATIONS OF LIABILITY IN THIS SECTION L.3 (LIMITS ON LIABILITY) APPLY: (I) TO THE MAXIMUM EXTENT PERMITTED BY APPLICABLE LAW; (II) TO LIABILITY IN TORT, INCLUDING FOR NEGLIGENCE; (III) REGARDLESS OF THE FORM OF ACTION, WHETHER IN CONTRACT, TORT, STRICT PRODUCT LIABILITY, OR OTHERWISE; (IV) EVEN IF THE BREACHING PARTY IS ADVISED IN ADVANCE OF THE POSSIBILITY OF THE DAMAGES IN QUESTION AND EVEN IF SUCH DAMAGES WERE FORESEEABLE; AND (E) EVEN IF THE INJURED PARTY'S REMEDIES FAIL OF

THEIR ESSENTIAL PURPOSE. THE FOREGOING WILL NOT APPLY TO (1) PERSONAL INJURY OR DEATH RESULTING FROM ANTHROPIC'S GROSS NEGLIGENCE OR WILLFUL MISCONDUCT; (2) FOR FRAUD; OR (3) FOR ANY OTHER MATTER FOR WHICH LIABILITY CANNOT BE EXCLUDED BY LAW.

d. The parties agree that they have entered into these Terms in reliance on the terms of this Section L.3 (Limits on Liability) and those terms form an essential basis of the bargain between the parties.

M. **Miscellaneous.**

1. **Notices.** All notices, demands, waivers, and other communications under these Terms (each, a "**Notice**") must be in writing. Except for notices related to demands to arbitrate or where equitable relief is sought, any Notices provided under these Terms may be delivered electronically to the address provided to Anthropic if to Customer; and to notices@anthropic.com if to Anthropic. Notice is effective only: (a) upon receipt by the receiving party, and (b) if the party giving the Notice has complied with all requirements of this Section M.1 (Notices).

2. **Electronic Communications.** Customer agrees to receive electronic communications from Anthropic based on Customer's use of the Services and related to these Terms. Except where prohibited by applicable law, electronic communications may be sent via email, through the Services or Customer's management dashboard, or posted on Anthropic's website. Anthropic may also provide electronic communications via text or SMS about Customer's use of the Services or as Customer otherwise requests from Anthropic. If Customer wishes to stop receiving such messages, Customer may request it from Anthropic or respond to any such texts with "STOP".

3. **Amendment and Modification.**

a. Anthropic may update the Usage Policy by providing email notice to Customer, but the update will not take effect for the then-current Order term unless Customer purchases additional Services during the then current-Order term. Such updates to the Usage Policy are intended to help Anthropic's customers and users stay safe and ensure Anthropics's Services are being used responsibility. If Customer believes the update is a material change as defined in GSAR Clause 552.212-4(w)(1)(vi), the parties agree to negotiate and incorporate mutually agreed upon changes to the Usage Policy.

b. Anthropic may update the Service Specific Terms which will apply (i) to each new Order, which take effect at the time the Order is executed; or (ii) the then-current Order term, but only as to updates that address terms and conditions to new features, services or products that are not Order Services under such Order. If Customer uses such new features, services or products during the then-current Order term, the new Service Specific Terms will apply.

c. If Anthropic intends to or does remove any countries from the Supported Regions Policy and Customer is using the Services in that country, Anthropic will provide Customer with email notice of such change and, except to the extent prohibited by applicable law, Customer will have six (6) months from the notice date to transition away from using the Services in any countries or regions removed from the Supported Regions Policy by Anthropic. Such updates to the Supported Regions Policy will not amend these Terms.

d. Changes to the Usage Policy, Service Specific Terms, or Support Regions Policy to the extent required to comply with changes to applicable laws or regulations will take effect immediately upon posting or Notice.

e. No other amendment to or modification of these Terms is effective unless it is in writing and signed by both parties. Failure to exercise or delay in exercising any rights or remedies arising from these Terms does not and will not be construed as a waiver; and no single or partial exercise of any right or remedy will preclude future exercise of such right or remedy.

f. Customer agrees that the updates set forth in (a) - (d) above are non-material changes in accordance with GSAR Clause 552.212-4(w)(1)(vi).

g. Changes will not apply retroactively. No other amendment to or modification of these Terms is effective unless it is in writing and signed by both parties. Failure to exercise or delay in exercising any rights or remedies arising from these Terms does not and will not be construed as a waiver; and no single or partial exercise of any right or remedy will preclude future exercise of such right or remedy.

4. **Federal Government End Use.** The Services constitute "commercial items" as defined in FAR 2.101. U.S. Government rights to use, modify, reproduce, release, perform, display, or disclose the Services, including any software, documentation, and technical data, are governed by these Terms in accordance with FAR 12.211- 12.212 and DFARS 227.7102-4 and 227.7202-4. For Department of Defense users, technical data customarily provided to the public will be furnished in accordance with DFARS 252.227-7015. Government agencies requiring additional rights must negotiate a written addendum to these Terms specifically granting those rights.

5. **Assignment and Delegation.** Neither party may assign its rights or delegate its obligations under these Terms without the other party's prior written consent, except that Anthropic may assign its rights and delegate its obligations to an affiliate or as part of a sale of all or substantially all its business in accordance with the provisions set forth in FAR 42.1204. Any purported assignment or delegation is null and void except as permitted above. No permitted assignment or delegation will relieve the contracting party or assignees of their obligations under these Terms. These Terms will bind and inure to the benefit of the parties and their respective permitted successors and assigns.

6. **Severability.** If a provision of these Terms is invalid, illegal, or unenforceable in any jurisdiction, such invalidity, illegality, or unenforceability will neither affect any other term or provision of these Terms nor invalidate or render unenforceable such term or provision in any other jurisdiction. Upon such determination that any term or other provision is invalid, illegal, or unenforceable, the parties will negotiate in good faith to modify these Terms to reflect the parties' original intent as closely as possible.

7. **Interpretation.** These Terms will be construed mutually, with neither party considered the drafter. Document and section titles are provided for convenience and will not be interpreted. The phrases "for example" or "including" or "or" are not limiting.

8. **Governing Law; Venue.** These Terms are governed by and construed in accordance with the Federal laws of the United States if Customer is an instrumentality of the U.S. Federal Government.

9. **Export and Sanctions.** Customer may not export or provide access to the Services to persons or entities or into countries or for uses where it is prohibited under U.S. or other applicable international law. Without limiting the foregoing sentence, this restriction applies (a) to countries where export from the US or into such country would be prohibited or illegal without first obtaining the appropriate license, and (b) to persons, entities, or countries covered by U.S. sanctions.

10. **Integration.** These Terms (including the Usage Policy, Supported Regions Policy, Service Specific Terms, DPA, Model Pricing Page and other documents or terms that are effective as of the Effective Date of these Terms) constitute the parties' entire understanding as to the Services' provision and use. These Terms supersede all other understandings or agreements between the parties regarding the Services. Any use of services or products offered by Anthropic that are not covered by an applicable Order are subject to Anthropic's standard Commercial Terms of Service. For clarity, the Model Pricing Page will not apply to these Terms.

11. **Force Majeure.** In accordance with GSAR Clause 552.212-4(f), neither party will be liable for failure or delay in performance to the extent caused by circumstances beyond its reasonable control.

This Anthropic Commercial Terms of Services is signed by the parties' authorized representatives on the dates below.

CUSTOMER	**ANTHROPIC, PBC**
By:	By:
Name:	Name:
Title:	Title:
Date:	Date:
Notices Email:	Notices Email: notices@anthropic.com

Service Specific Terms

Effective May 1, 2025

These Service Specific Terms ("Service Specific Terms") contain supplemental terms that apply to specific Services and form part of the agreement you or the organization, company, or other entity that you represent ("Customer") have with Anthropic for such Services that links to or otherwise references these Service Specific Terms, such as the Commercial Terms of Service ("Agreement"). Any capitalized terms not defined in these Service Specific Terms will have the definitions stated in the Agreement. In the event of any conflict between the Terms and these Service Specific Terms, these Service Specific Terms control for the relevant Services.

A. Claude for Work (Team Plan; Enterprise Plan)

Customer must inform its Users that (i) they are accessing an administered service offering that enables Customer's access to and control over data submitted to the Services by Customer and its Users, and (ii) use of the Services are subject to Anthropic's policies, including the Privacy Policy and Usage Policy. Customer is responsible for obtaining any necessary User consents and providing all User notices required by applicable law.

B. Beta Services

Anthropic may offer Services that are in pre-release, beta, preview or similar trial form ("Beta Services"). This means that they are not suitable for production use and provided "*as-is*" on a temporary basis. Anthropic is not responsible for Customer's use of or reliance on Beta Services and is not obligated to provide any indemnity for their use. Anthropic's liability related to Beta Services is limited to the lesser of $1,000 and Fees paid by Customer for the Services in the previous 12 months.

C. Fine-Tuning Services

Anthropic may offer Services that allow Customer to fine-tune or conduct other forms of model customization ("Fine-Tuning Services" and any such resulting model, a "Fine-Tuned Model"), including using data provided by Customer ("Fine-Tuning Materials"). Fine-Tuning Services use the Fine-Tuning Materials for the purpose of generating or validating Fine-Tuned Models. Fine-Tuning Materials are Customer Content. Customer's use of the Fine-Tuning Services does not constitute a violation of the prohibition on training in the

Agreement; provided that the training conducted is solely for Customer's use of the Fine-Tuning Services.

D. Cloud Hosted Services

1. Generally. Anthropic offers certain Services through the following third-party managed and hosted technical environments: (1) Amazon Bedrock ("Bedrock"), which is hosted and managed by Amazon Web Services, Inc. and affiliates ("AWS"); and (2) Vertex AI ("Vertex"), which is hosted and managed by Google LLC and affiliates ("Google"). AWS and Google are "Cloud Providers" and their environments are "Cloud Platforms" for purposes of these Services Specific Terms. Customer acknowledges and agrees that Anthropic is not responsible for Cloud Platforms or any other services Customer receives from Cloud Providers, and that changes to Cloud Platforms or other Cloud Provider services may materially impact Customer's use of the Services. Customer may only use the Services in compliance with applicable Cloud Provider policies and agreements it has with the applicable Cloud Provider.
2. Content on Cloud Platforms. The technology provided by Anthropic to enable access to the Services via the Cloud Platforms does not give Anthropic access to any Customer instance with the Cloud Providers, including Customer Content they contain.
3. Fees. Customer must pay all applicable fees related to use of the Services via Cloud Platforms to the applicable Cloud Provider at Anthropic's then-current rates in accordance with the applicable GSA Schedule Pricelist, unless otherwise agreed by the parties.

E. Development Partner Program

Customer may elect (in its sole discretion) to participate in the Development Partner Program by enabling the permissive data opt-in setting for the Services ("Development Partner Mode"). If Customer enables Development Partner Mode, Anthropic may use the data that Customer submits to the Services (e.g., Customer Content) in connection with Anthropic's products and services, including to train models.

Privacy Policy

Effective May 1, 2025

Anthropic is an AI safety and research company working to build reliable, interpretable, and steerable AI systems.

This Privacy Policy explains how we collect, use, disclose, and process your personal data when you use our website and other places where Anthropic acts as a *data controller*—for example, when you interact with Claude.ai or other products as a consumer for personal use ("Services") or **when Anthropic operates and provides our commercial customers and their end users with access to our commercial products, such as the Claude Team plan (“**Commercial Services**”).**

This Privacy Policy does not apply where Anthropic acts as a *data processor* and processes personal data on behalf of commercial customers using Anthropic’s Commercial Services – for example, your employer has provisioned you a Claude for Work account, or you're using an app that is powered on the back-end with Claude. In those cases, the commercial customer is the controller, and you can review their policies for more information about how they handle your personal data.

For information about how we collect and use personal data to develop our language models that power our Services, when we may receive and process personal data of non-users, and your choices with respect to that information, please see our Notice on Model Training.

This Privacy Policy also describes your privacy rights. More information about your rights, and how to exercise them, is set out in Section 4 (“Rights and Choices”).

If you are located in Canada, please read section 11 of the Privacy Policy which applies to you.

If you are located in Brazil, please read section 12 of the Privacy Policy which applies to you.

1. Collection of Personal Data

We collect the following categories of personal data:

Personal data you provide to us directly

- Identity and Contact Data: Anthropic collects identifiers, including your name, email address, and phone number when you sign up for an Anthropic account, or to receive information on our Services. We may also collect or generate indirect identifiers (e.g., “USER12345”).

- Payment Information: We shall collect your payment information if you choose to purchase access to Anthropic’s products and services.

- Inputs and Outputs: Our AI services allow you to interact with the Services in a variety of formats (“Prompts” or "Inputs"), which generate responses (“Outputs”) based on your Inputs. This includes where you choose to integrate third-party applications with our services. If you include personal data or reference external content in your Inputs, we will collect that information and this information may be reproduced in your Outputs.

- Feedback on your use of our Services: We appreciate feedback, including ideas and suggestions for improvement or rating an Output in response to an Input ("Feedback"). If you rate an Output in response to an Input—for example, by using the thumbs up/thumbs down icon—we will store the entire related conversation as part of your Feedback. You can learn more about how we use Feedback here.

- Communication Information: If you communicate with us, including via our chatbot on our Help site, we collect your name, contact information, and the contents of any messages you send.

Personal data we receive automatically from your use of the Services

When you use the Services, we also receive certain technical data automatically (described below, collectively "Technical Information"). This includes:

- Device and Connection Information. Consistent with your device or browser permissions, your device or browser automatically sends us information about when and how you install, access, or use our Services. This includes information such as your device type, operating system information, browser information and web page referers, mobile network, connection information, mobile operator or internet service provider (ISP), time zone setting, IP address (including information about the location of the device derived from your IP address), identifiers (including device or advertising identifiers, probabilistic identifiers, and other unique personal or online identifiers).

- Usage Information. We collect information about your use of the Services, such as the dates and times of access, browsing history, search, information about the links you click, pages you view, and other information about how you use the Services, and technology on the devices you use to access the Services.

- Log and Troubleshooting Information. We collect information about how our Services are performing when you use them. This information includes log files. If you or your device experiences an error, we may collect information about the error, the time the error occurred, the feature being used, the state of the application when the error occurred, and any communications or content provided at the time the error occurred.

- Cookies & Similar Technologies. We and our service providers use cookies, scripts, or similar technologies ("Cookies") to manage the Services and to collect information about you and your use of the Services. These technologies help us to recognize you, customize or personalize your experience, market additional products or services to you, and analyze the use of our Services to make them safer and more

useful to you. For more details about how we use these technologies, and your opt-out controls and other options, please visit our Cookie Policy.

Personal data we collect or receive to train our models

Anthropic obtains personal data from third party sources in order to train our models. Specifically, we train our models using data from the following sources:

- Publicly available information via the Internet
- Datasets that we obtain through commercial agreements with third party businesses
- Data that our users or crowd workers provide
- Data that we generate internally

For more information about how we collect and use personal data to develop our language models that power our Services, the steps we take to minimize the privacy impact on individuals through the training process, and your choices with respect to that information, please see our separate Notice on Model Training.

2. Uses of Personal Data Permitted Under Applicable Data Protection Laws

We use your personal data for the following purposes:

- To provide, maintain and facilitate any products and services offered to you with respect to your Anthropic account, which are governed by our Terms of Service;
- To provide, maintain and facilitate optional services and features that enhance platform functionality and user experience;

- To communicate with you, including to send you information about our Services and events;
- To create and administer your Anthropic account;
- To facilitate payments for products and services provided by Anthropic;
- To prevent and investigate fraud, abuse, and violations of our Usage Policy, unlawful or criminal activity, unauthorized access to or use of personal data or Anthropic systems and networks, to protect our rights and the rights of others, and to meet legal, governmental and institutional policy obligations;
- To investigate and resolve disputes;
- To investigate and resolve security issues;
- To debug and to identify and repair errors that impair existing functionality
- To improve the Services and conduct research; and
- To enforce our Terms of Service and similar terms and agreements, including our Usage Policy.

We will not use your Inputs or Outputs to train our models, unless: (1) your conversations are flagged for Trust & Safety review (in which case we may use or analyze them to improve our ability to detect and enforce our Usage Policy, including training models for use by our Trust and Safety team, consistent with Anthropic's safety mission), or (2) you've explicitly reported the materials to us (for example via our feedback mechanisms), or (3) you've otherwise explicitly opted in to the use of your Inputs and Outputs for training purposes.

Please see Section 10 below for details of our legal bases for processing your personal data.

3. How We Disclose Personal Data

Anthropic will disclose personal data to the following categories of third parties for the purposes explained in this Policy:

- Affiliates & corporate partners. Anthropic discloses the categories of personal data described above between and among its affiliates and related entities.

- Service providers & business partners. Anthropic may disclose the categories of personal data described above with service providers and business partners for a variety of business purposes, including website and data hosting, ensuring compliance with industry standards, research, auditing, data processing, and providing you with the services.

Anthropic may also disclose personal data in the following circumstances:

- As part of a significant corporate event. If Anthropic is involved in a merger, corporate transaction, bankruptcy, or other situation involving the transfer of business assets, Anthropic will disclose your personal data as part of these corporate transactions.

- Third-Party Websites and Services: Our Services may involve integrations with, or may direct you to, websites, apps, and services managed by third parties. By interacting with these third parties, you are providing information directly to the third party and not Anthropic and subject to the third party's privacy policy.If you access third-party services, such as social media sites or other sites linked through the Services (e.g., if you follow a link to our Twitter account), these third-party services will be able to collect personal data about you, including information about your activity on the Services. If we link to a site or service via our Services, you should read their data usage policies or other documentation. Our linking to another site or service doesn't mean we endorse it or speak for that third party.

- Pursuant to regulatory or legal requirements, safety, rights of others, and to enforce our rights or our terms. We may disclose personal data to governmental regulatory authorities as required by law, including for legal, tax or accounting purposes, in response to their requests for such information or to assist in investigations. We may also disclose personal data to third parties in connection with claims, disputes or litigation, when otherwise permitted or required by law, or if we determine its disclosure is necessary to protect the health and safety of you or any other person, to protect against fraud or credit risk, to enforce our legal rights or the legal rights of others, to enforce contractual commitments that you have made, or as otherwise permitted or required by applicable law.

- With an individual's consent. Anthropic will otherwise disclose personal data when an individual gives us permission or directs us to disclose this information, including as a part of our Services.

You can find information on our Subprocessor List about the third parties Anthropic engages to help us process personal data provided to us where Anthropic acts as a data processor, such as with respect to personal data we receive, process, store, or host when you use Anthropic's commercial services.

4. Rights and Choices

Depending on where you live and the laws that apply in your country of residence, you may enjoy certain rights regarding your personal data, as described further below. However, please be aware that these rights are limited, and that the process by which we may need to action your requests regarding our training dataset are complex. We may also decline a request if we have a lawful reason for doing so. That said, we strive to prioritize the protection of personal data, and comply with all applicable privacy laws.

To exercise your rights, you or an authorized agent may submit a request by emailing us at privacy@anthropic.com. After we receive your request, we may verify it by requesting

information sufficient to confirm your identity. You may also have the right to appeal requests that we deny by emailing privacy@anthropic.com. Anthropic will not discriminate based on the exercising of privacy rights you may have. Set out below is a summary of the rights which you may enjoy, depending on the laws that apply in your country of residence.

- Right to know: the right to know what personal data Anthropic processes about you, including the categories of personal data, the categories of sources from which it is collected, the business or commercial purposes for collection, and the categories of third parties to whom we disclose it.

- Access & data portability: the right to request a copy of the personal data Anthropic processes about you, subject to certain exceptions and conditions. In certain cases and subject to applicable law, you have the right to port your information.

- Deletion:the right to request that we delete personal data collected from you when you use our Services, subject to certain exceptions. You also are able to delete individual conversations, which will be removed immediately from your conversation history and automatically deleted from our back-end within 30 days. Learn more here.

- Correction:the right to request that we correct inaccurate personal data Anthropic retains about you, subject to certain exceptions. Please note that we cannot guarantee the factual accuracy of Outputs. If Outputs contain factually inaccurate personal data relating to you, you can submit a correction request and we will make a reasonable effort to correct this information—but due to the technical complexity of our large language models, it may not always be possible for us to do so.

- Objection: the right to object to processing of your personal data, including profiling conducted on grounds of public or legitimate interest. In places where such a right applies, we will no longer process the personal data in case of such objection unless we demonstrate compelling legitimate grounds for the processing which override

your interests, rights, and freedoms, or for the establishment, exercise or defense of legal claims. If we use your information for direct marketing, you can object and opt out of future direct marketing messages using the unsubscribe link in such communications.

- Restriction: the right to restrict our processing of your personal data in certain circumstances.

- Withdrawal of consent. Where Anthropic's processing of your personal data is based on consent, you have the right to withdraw your consent. The withdrawal of consent will not affect the lawfulness of processing based on consent before its withdrawal.

- Automated decision-making: Anthropic does not engage in decision making based solely on automated processing or profiling in a manner which produces a legal effect (i.e., impacts your legal rights) or significantly affects you in a similar way (e.g., significantly affects your financial circumstances or ability to access essential goods or services).

- Sale & targeted Anthropic marketing of its products and services. Anthropic does not "sell" your personal data as that term is defined by applicable laws and regulations. You can opt-out of sharing your personal data for targeted advertising to promote our products and services, and we will honor global privacy controls. To learn more, click here.

5. Data Transfers

When you access our website or Services, your personal data may be transferred to our servers in the US, or to other countries outside the European Economic Area ("EEA") and the UK. This may be a direct provision of your personal data to us, or a transfer that we or a third party make.

Where information is transferred outside the EEA or the UK, we ensure it benefits from an adequate level of data protection by relying on:

- Adequacy decisions. These are decisions from the European Commission under Article 45 GDPR (or equivalent decisions under other laws) where they recognise that a country outside of the EEA offers an adequate level of data protection. We transfer your information as described in "Collection of Personal Data" to some countries with adequacy decisions, such as the countries listed here; or

- Standard contractual clauses. The European Commission has approved contractual clauses under Article 46 GDPR that allows companies in the EEA to transfer data outside the EEA. These (and their approved equivalent for the UK and Switzerland) are called standard contractual clauses. We rely on standard contractual clauses to transfer information as described in "Collection of Personal Data" to certain affiliates and third parties in countries without an adequacy decision.

In certain situations, we rely on derogations provided for under applicable data protection law to transfer information to a third country.

6. Data Retention, Data Lifecycle, and Security Controls

Anthropic retains your personal data for as long as reasonably necessary for the purposes and criteria outlined in this Privacy Policy and explained further in our privacy center.

When the personal data collected is no longer required by us, we and our service providers will perform the necessary procedures for destroying, deleting, erasing, or converting it into an anonymous form as permitted or required under applicable laws.

Aggregated or De-Identified Information

We may process personal data in an aggregated or de-identified form to analyze the effectiveness of our Services, conduct research, study user behavior, and train our AI models as permitted under applicable laws. For instance:

- When you submit Feedback, we disassociate Inputs and Outputs from your user ID to use them for training and improving our models.
- If our systems flag Inputs or Outputs for potentially violating our Usage Policy, we disassociate the content from your user ID to train our trust and safety internal classification and generative models. However, we may re-identify the Inputs or Outputs to enforce our Usage Policy with the responsible user if necessary.
- To improve user experience, we may analyze and aggregate general user behavior and usage data. This information does not identify individual users.

Security Controls Relating to our Processing of Personal Data

We implement appropriate technical and organizational security measures designed to protect personal data from loss, misuse, and unauthorized access, disclosure, alteration, or destruction.

7. Children

Our Services are not directed towards, and we do not knowingly collect, use, disclose, sell, or share any information about, children under the age of 18. If you become aware that a child under the age of 18 has provided any personal data to us while using our Services, please email us at privacy@anthropic.com and we will investigate the matter and, if appropriate, delete the personal data.

8. Changes to Our Privacy Policy

Anthropic may update this Privacy Policy from time to time. We will notify you of any material changes to this Privacy Policy, as appropriate, and update the Effective Date at the top of https://www.anthropic.com/legal/privacy. You can view a summary of privacy policy changes and previous versions in our Privacy Center.

9. Contact Information

If you live in the European Economic Area (EEA), UK or Switzerland (the "European Region"), the data controller responsible for your personal data is Anthropic Ireland, Limited. If you live outside the European Region, the data controller responsible for your personal data is Anthropic PBC.

If you have any questions about this Privacy Policy, or have any questions, complaints or requests regarding your personal data, you can contact us as described below:

- Anthropic PBC with a registered address at 548 Market St, PMB 90375, San Francisco, CA 94104 (United States).
- Anthropic Ireland, Limited with a registered address at 6th Floor, South Bank House, Barrow Street. Dublin 4, D04 TR29 (Ireland).

You can email us at privacy@anthropic.com and contact our Data Protection Officer at dpo@anthropic.com.

Please note that under many countries' laws, you have the right to lodge a complaint with the supervisory authority in the place in which you live or work. A full list of EU supervisory authorities' contact details is available here. If you live or work in the UK, you have the right to lodge a complaint with the UK Information Commissioner's Office. If you live in Brazil, you have the right to lodge a complaint with the Brazilian Data Protection Authority (ANPD).If you

live in Australia, you have the right to lodge a complaint with the Office of the Australian Information Commissioner.

10. Legal Bases for Processing

Purpose	Type of Data	Legal Basis
To provide, maintain and facilitate any products and services offered to you with respect to your Anthropic account, which are governed by our Terms of Service	• Identity and Contact Data • Payment Information • Feedback • Inputs and Outputs	• Contract
To provide, maintain and facilitate optional services and features that enhance platform functionality and user experience	• Identity and Contact Data • Feedback • Inputs and Outputs	• Legitimate interests It is in our and our users' legitimate interests to expand our product features and deliver additional services that enhance platform functionality and user experience.
To communicate with you, including to send you information about our Services and events	• Identity and Contact Data	• Where necessary to perform a contract with you, such as processing

	• Communication Information • Technical Information	your contact information to send you a technical announcement about the Services. • Your consent when we ask for it to process your personal data for a specific purpose that we communicate to you, such as processing your contact information to send you certain forms of marketing communications.
To create and administer your Anthropic account	• Identity and Contact Data • Payment Information • Feedback	• Contract
To facilitate payments for products and services provided by Anthropic	• Identity and Contact Data • Payment Information	• Contract

To prevent and investigate fraud, abuse, and violations of our Usage Policy, unlawful or criminal activity, unauthorized access to or use of personal data or Anthropic systems and networks, to protect our rights and the rights of others, and to meet legal, governmental and institutional policy obligations	• Identity and Contact Data • Payment Information • Inputs and Outputs • Technical Information	• Legitimate interests • Legal obligation It is our legitimate interest to protect our business, employees and users from illegal activities, inappropriate behavior or violations of terms that would be detrimental. We also have a duty to cooperate with authorities.
To investigate and resolve disputes	• Identity and Contact Data • Inputs and Outputs • Feedback	• Legitimate interests • Legal obligation It is our legitimate interest to fully understand and make reasonable efforts to resolve customer complaints in order to improve user satisfaction. We also have a legal obligation in some cases.
To investigate and resolve security issues	• Identity and Contact Data • Feedback • Technical Information	• Legal obligation • Legitimate interests It is our legitimate interest to protect user data and our systems from intrusion or compromise through monitoring and swift response.

	• Inputs and Outputs	We also have a legal obligation to provide adequate security safeguards.
To debug and to identify and repair errors that impair existing functionality	• Identity and Contact Data • Feedback • Technical Information	• Legitimate interests It is our legitimate interest to maintain continuous functioning of our services and rapid correction of problems to ensure a positive user experience that encourages engagement.
To improve the Services and conduct research (excluding model training)	• Identity and Contact Data • Feedback • Technical Information • Inputs and Outputs	• Legitimate interests It is our legitimate interest and in the interest of Anthropic users to evaluate the use of the Services and adoption of new features to inform the development of future features and improve direction and development of the Services.
To improve the Services and conduct research (including model training). See our Notice on Model Training for more details on the data used to train our models.	• Feedback (which may include Inputs and Outputs)	• Consent (when users submit Feedback)

	• Inputs and Outputs flagged for Trust & Safety review • Data provided through the Development Partner Program	• Legitimate interests (for trust and safety purposes) It is our legitimate interest and in the interest of Anthropic users to evaluate the use of the Services and adoption of new features to inform the development of future features and improve direction and development of the Services.
To enforce our Terms of Service and similar terms and agreements, including our Usage Policy.	• Identity and Contact Data • Inputs and Outputs • Technical Information	• Contract • Legitimate interests In certain circumstances outside of the performance of our contract with you, we may rely on legitimate interests.It is our legitimate interest to enforce the rules and policies governing use of our services, to maintain intended functionality and value for users. We aim to provide a safe, useful platform.

Use of personal data

11. Supplemental Disclosures for Residents of Canada

These supplemental disclosures contain additional information relevant to residents of Canada. This content should be read in conjunction with the rest of our Privacy Policy. In case of conflict

between our Privacy Policy and these supplemental disclosures, the supplemental disclosures shall prevail in relation to residents of Canada.

Consent. By expressly consenting to this Privacy Policy, you confirm you have read, understand, and consent to the collection, use, processing, and disclosure of your personal data in accordance with this Privacy Policy and understand that, in jurisdictions where it is available, Anthropic also relies on other lawful bases for the foregoing as more fully set out in this policy. We will only collect, use and disclose your personal data with your consent, unless otherwise permitted or required by law. Your consent may be given expressly or implied, depending on the circumstances and the sensitivity of the information involved. You may withdraw consent at any time, subject to legal or contractual restrictions and reasonable notice.

Cross-jurisdictional Transfers. By providing us with personal data, you acknowledge and agree that your personal data may be transferred or disclosed to other jurisdictions for processing and storage outside of Canada, including to the United States and the countries listed on our Subprocessor List, where laws regarding the protection of personal data may be less stringent than the laws in your jurisdiction. Furthermore, we may disclose your personal data in these jurisdictions in response to legal processes or where we believe in good faith that disclosure is required or permitted by law.

Contact. If you have any questions or comments about our processing of your personal data, or to exercise your rights as outlined in Section 4. ("Rights and Choices"), please contact us at privacy@anthropic.com.

12. Supplemental Disclosures for Residents of Brazil

These supplemental disclosures contain additional information relevant to residents of Brazil. This content should be read in conjunction with the rest of our Privacy Policy. In case of conflict between our Privacy Policy and these supplemental disclosures, the supplemental disclosures shall prevail in relation to residents of Brazil.

Legal Bases. Depending on the specific purpose of the processing, we may rely on different grounds than those listed under section 2, where permitted by and in accordance with the Brazilian General Data Protection Law (LGPD). For example, we may rely on the "exercise of legal rights" basis to process personal data associated with customer complaints and to enforce our Terms of Service and similar terms and agreements, including our Usage Policy.

Data Subject's Rights. LGPD grants certain rights regarding your personal data, which differ from the ones listed under section 4. We will respond to your requests to exercise your rights below in accordance with applicable law:

- Confirmation of whether your data is being processed. You have the right to receive a confirmation on whether Anthropic processes your data.Access to your data. You have the right to know what personal data Anthropic processes about you.

- Correction of incomplete, inaccurate or outdated data. You have the right to request the correction of your data that is incomplete, inaccurate, or outdated.

- Anonymization, blocking or erasure of data. You have the right to request the anonymisation, blocking or erasure of data that is unnecessary, excessive or processed in non-compliance with the provisions of the law.

- Portability of personal data to a third party. You have the right to request portability of your data to a third-party, as long as this does not infringe on our trade secrets.

- Information of public and private entities with which we shared data. You have the right to request information of public and private entities with which we have shared your data.

- Information about the possibility to refuse to provide consent and the respective consequences, when applicable.

- Withdrawal of your consent. You have the right to withdraw your consent. This procedure will be carried out free of charge.

- Request a review of decisions made solely based on automated processing of personal data.

Please keep in mind that these rights are not absolute and may not apply in certain circumstances. For example, in certain cases we may continue to process and retain data regardless of your request for deletion, objection, blocking or anonymisation, in order to comply with legal, contractual and regulatory obligations, safeguard and exercise rights, including in judicial, administrative and arbitration proceedings and in other cases provided for by law.

International Data Transfers. You acknowledge that Anthropic is a company based and headquartered in the United States. Any information we hold about you will be transferred to, used, processed, and stored in the United States and other countries and territories, which may not have data privacy or data protection laws equivalent to those in your country or territory. For the proper operation of the Services, Anthropic needs to carry out international transfers of personal data. In the case of Brazil, we will rely on standard contractual clauses (SCCs) for our data transfers where required and in instances where they are not covered by an adequacy decision. These SCCs have been approved by the Brazilian Data Protection Authority (ANPD), which is the "competent supervisory authority" for these transfers, as governed by Brazilian Data Protection Laws. You can view the SCCs adopted by the ANPD here.

Usage Policy

Effective June 6, 2024

Our Usage Policy (also referred to as our "Acceptable Use Policy" or "AUP") applies to anyone who uses Anthropic's products and services, and is intended to help our users stay safe and ensure our products and services are being used responsibly.

The Usage Policy is categorized according to who can use our products and for what purposes. We will update our policy as our technology and the associated risks evolve or as we learn about unanticipated risks from our users.

- Universal Usage Standards: Our Universal Usage Standards apply to all users including individuals, developers, and businesses.
- High-Risk Use Case Requirements: Our High-Risk Use Case Requirements apply to specific use cases that pose an elevated risk of harm.
- Disclosure Requirements: Our Disclosure Requirements apply to specific use cases where it is especially important for users to understand that they are interacting with an AI system.

Anthropic's Trust and Safety Team will implement detections and monitoring to enforce our Usage Policies so please review these policies carefully before using our products. If we learn that you have violated our Usage Policy, we may throttle, suspend, or terminate your access to our products and services. If you discover that our model outputs are inaccurate, biased or harmful, please notify us at usersafety@anthropic.com or report it directly in the product through the "report issues" thumbs down button. You can read more about our Trust and Safety practices and recommendations in our T&S Support Center.

This Usage Policy is calibrated to strike an optimal balance between enabling beneficial uses and mitigating potential harms. Anthropic may enter into contracts with certain governmental customers that tailor use restrictions to that customer's public mission and legal authorities if, in Anthropic's judgment, the contractual use restrictions and applicable safeguards are adequate to mitigate the potential harms addressed by this Usage Policy.

Universal Usage Standards

Do Not Compromise Children's Safety

This includes using our products or services to:

- Create, distribute, or promote child sexual abuse material. We strictly prohibit and will report to relevant authorities and organizations where appropriate any content that exploits or abuses minors
- Facilitate the trafficking, sextortion, or any other form of exploitation of a minor
- Facilitate minor grooming, including generating content designed to impersonate a minor
- Facilitate or depict child abuse of any form, including instructions for how to conceal abuse
- Promote or facilitate pedophilic relationships, including via roleplay with the model
- Fetishize minors

Do Not Compromise Critical Infrastructure

This includes using our products or services to:

- Facilitate the destruction or disruption of critical infrastructure such as power grids, water treatment facilities, telecommunication networks, or air traffic control systems
- Obtain unauthorized access to critical systems such as voting machines, healthcare databases, and financial markets
- Interfere with the operation of military bases and related infrastructure

Do Not Incite Violence or Hateful Behavior

This includes using our products or services to:

- Incite, facilitate, or promote violent extremism, terrorism, or hateful behavior
- Depict support for organizations or individuals associated with violent extremism, terrorism, or hateful behavior
- Facilitate or promote any act of violence or intimidation targeting individuals, groups, animals, or property
- Promote discriminatory practices or behaviors against individuals or groups on the basis of one or more protected attributes such as race, ethnicity, religion, nationality, gender, sexual orientation, or any other identifying trait

Do Not Compromise Someone's Privacy or Identity

This includes using our products or services to:

- Compromise security or gain unauthorized access to computer systems or networks, including spoofing and social engineering
- Violate the security, integrity, or availability of any user, network, computer, device, or communications system, software application, or network or computing device
- Violate any person's privacy rights as defined by applicable privacy laws, such as sharing personal information without consent, accessing private data unlawfully, or violating any relevant privacy regulations
- Misuse, collect, solicit, or gain access to private information without permission such as non-public contact details, health data, biometric or neural data (including facial recognition), or confidential or proprietary data
- Impersonate a human by presenting results as human-generated, or using results in a manner intended to convince a natural person that they are communicating with a natural person when they are not

Do Not Create or Facilitate the Exchange of Illegal or Highly Regulated Weapons or Goods

This includes using our products or services to:

- Produce, modify, design, market, or distribute weapons, explosives, dangerous materials or other systems designed to cause harm to or loss of human life
- Engage in or facilitate any illegal activity, such as the use, acquisition, or exchange of illegal and controlled substances, or the facilitation of human trafficking and prostitution

Do Not Create Psychologically or Emotionally Harmful Content

This includes using our products or services to:

- Facilitate or conceal any form of self-harm, including disordered eating and unhealthy or compulsive exercise
- Engage in behaviors that promote unhealthy or unattainable body image or beauty standards, such as using the model to critique anyone’s body shape or size
- Shame, humiliate, intimidate, bully, harass, or celebrate the suffering of individuals
- Coordinate the harassment or intimidation of an individual or group
- Generate content depicting sexual violence

- Generate content depicting animal cruelty or abuse
- Generate violent or gory content that is inspired by real acts of violence
- Promote, trivialize, or depict graphic violence or gratuitous gore
- Develop a product, or support an existing service that facilitates deceptive techniques with the intent of causing emotional harm

Do Not Spread Misinformation

This includes the usage of our products or services to:

- Create and disseminate deceptive or misleading information about a group, entity or person
- Create and disseminate deceptive or misleading information about laws, regulations, procedures, practices, standards established by an institution, entity or governing body
- Create and disseminate deceptive or misleading information with the intention of targeting specific groups or persons with the misleading content
- Create and advance conspiratorial narratives meant to target a specific group, individual or entity
- Impersonate real entities or create fake personas to falsely attribute content or mislead others about its origin without consent or legal right
- Provide false or misleading information related to medical, health or science issues

Do Not Create Political Campaigns or Interfere in Elections

This includes the usage of our products or services to:

- Promote or advocate for a particular political candidate, party, issue or position. This includes soliciting votes, financial contributions, or public support for a political entity
- Engage in political lobbying to actively influence the decisions of government officials, legislators, or regulatory agencies on legislative, regulatory, or policy matters. This includes advocacy or direct communication with officials or campaigns to sway public opinion on specific legislation or policies

- Engage in campaigns, including political campaigns, that promote false or misleading information to discredit or undermine individuals, groups, entities or institutions
- Incite, glorify or facilitate the disruption of electoral or civic processes, such as targeting voting machines, or obstructing the counting or certification of votes
- Generate false or misleading information on election laws, procedures and security, candidate information, how to participate, or discouraging participation in an election

Do Not Use for Criminal Justice, Law Enforcement, Censorship or Surveillance Purposes

This includes the usage of our products or services to:

- Make determinations on criminal justice applications, including making decisions about or determining eligibility for parole or sentencing
- Target or track a person's physical location, emotional state, or communication without their consent, including using our products for facial recognition, battlefield management applications or predictive policing
- Utilize Claude to assign scores or ratings to individuals based on an assessment of their trustworthiness or social behavior
- Build or support emotional recognition systems or techniques that are used to infer people's emotions
- Analyze or identify specific content to censor on behalf of a government organization
- Utilize Claude as part of any biometric categorization system for categorizing people based on their biometric data to infer their race, political opinions, trade union membership, religious or philosophical beliefs, sex life or sexual orientation
- Use the model for any official local, state or national law enforcement application. Except for the following permitted applications by law enforcement organizations:
    - Back office uses including internal training, call center support, document summarization, and accounting;
    - Analysis of data for the location of missing persons, including in human trafficking cases, and other related applications, provided that such

applications do not otherwise violate or impair the liberty, civil liberties, or human rights of natural persons

Do Not Engage in Fraudulent, Abusive, or Predatory Practices

This includes using our products or services to:

- Facilitate the production, acquisition, or distribution of counterfeit or illicitly acquired goods
- Promote or facilitate the generation or distribution of spam
- Generate content for fraudulent activities, schemes, scams, phishing, or malware that can result in direct financial or psychological harm
- Generate content for the purposes of developing or promoting the sale or distribution of fraudulent or deceptive products
- Generate deceptive or misleading digital content such as fake reviews, comments, or media
- Engage in or facilitate multi-level marketing, pyramid schemes, or other deceptive business models that use high-pressure sales tactics or exploit participants
- Promote or facilitate payday loans, title loans, or other high-interest, short-term lending practices that exploit vulnerable individuals
- Engage in deceptive, abusive behaviors, practices, or campaigns that exploits people due to their age, disability or a specific social or economic situation
- Promote or facilitate the use of abusive or harassing debt collection practices
- Develop a product, or support an existing service that deploys subliminal, manipulative, or deceptive techniques to distort behavior by impairing decision-making
- Plagiarize or engage in academic dishonesty

Do Not Abuse our Platform

This includes using our products or services to:

- Coordinate malicious activity across multiple accounts such as creating multiple accounts to avoid detection or circumvent product guardrails or generating identical or similar prompts that otherwise violate our Usage Policy
- Utilize automation in account creation or to engage in spammy behavior

- Circumvent a ban through the use of a different account, such as the creation of a new account, use of an existing account, or providing access to a person or entity that was previously banned
- Facilitate or provide account access to Claude to persons or entities who are located in unsupported locations
- Intentionally bypass capabilities or restrictions established within our products for the purposes of instructing the model to produce harmful outputs (e.g., jailbreaking or prompt injection) without an authorized use-case approved by Anthropic
- Unauthorized utilization of prompts and completions to train an AI model (e.g., "model scraping")

Do Not Generate Sexually Explicit Content

This includes the usage of our products or services to:

- Depict or request sexual intercourse or sex acts
- Generate content related to sexual fetishes or fantasies
- Facilitate, promote, or depict incest or bestiality
- Engage in erotic chats

High-Risk Use Case Requirements

Some integrations (meaning use cases involving the use of our products and services) pose an elevated risk of harm because they influence domains that are vital to public welfare and social equity. "High-Risk Use Cases" include:

- Legal: Integrations related to legal interpretation, legal guidance, or decisions with legal implications
- Healthcare: Integrations affecting healthcare decisions, medical diagnosis, patient care, or medical guidance. Wellness advice (e.g., advice on sleep, stress, nutrition, exercise, etc.) does not fall under this category
- Insurance: Integrations related to health, life, property, disability, or other types of insurance underwriting, claims processing, or coverage decisions
- Finance: Integrations related to financial decisions, including investment advice, loan approvals, and determining financial eligibility or creditworthiness

- Employment and housing: Integrations related to decisions about the employability of individuals, resume screening, hiring tools, or other employment determinations or decisions regarding eligibility for housing, including leases and home loans
- Academic testing, accreditation and admissions: Integrations related to standardized testing companies that administer school admissions (including evaluating, scoring or ranking prospective students), language proficiency, or professional certification exams; agencies that evaluate and certify educational institutions.
- Media or professional journalistic content: Integrations related to using our products or services to automatically generate content and publish it for external consumption

If your integration is listed above, we require that you implement the additional safety measures listed below:

- Human-in-the-loop: when using our products or services to provide advice, recommendations, or subjective decisions that directly impact individuals in high-risk domains, a qualified professional in that field must review the content or decision prior to dissemination or finalization. This requirement applies specifically to content or decisions that are provided to consumers or the general public, or decisions made about an individual. Your business is responsible for the accuracy and appropriateness of that information. For other types of content generation or interactions with users that do not involve direct advice, recommendations, or subjective decisions, human review is strongly encouraged but not mandatory.
- Disclosure: you must disclose to your customers or end users that you are using our services to help inform your decisions or recommendations.

Disclosure Requirements

Finally, the below use cases – regardless of whether they are High Risk Use Cases – must disclose to their users that they are interacting with an AI system rather than a human:

- All customer-facing chatbots including any external-facing or interactive AI agent
- Products serving minors: Organizations providing minors with the ability to directly interact with products that incorporate our API(s). Note: These organizations must also comply with the additional guidelines outlined in our Help Center article

Terms and Conditions for LLM Developers' Government Addendum

(a) *Applicability*

(1) These terms and conditions apply when Government Data will be processed by a Large Language Model Artificial Intelligence System (LLM).

(2) These terms and conditions do not apply when—

(i) The LLM is embedded in a common commercial product, such as a word processor or map navigation system (see Section 7223(4)(B) of Pub. L. 117-263);

(ii) The LLM functionality is incidental to the primary purpose of the core requirement being procured.

(b) *Definitions.* As used in these terms and conditions —

Background Data means any pre-existing proprietary content, reference materials, knowledge bases, or other intellectual property owned or controlled by the LLM Developer that may be referenced, retrieved, augmented, or otherwise incorporated into the LLM's processing or outputs through any enrichment mechanism, including but not limited to retrieval processes, vector stores, embeddings, knowledge graphs, plugins, tool calls, agent actions, or similar mechanisms.

Custom Development means any design of or modifications, customizations, configurations, or enhancements to LLMs or associated implementations or workflows, and any related work product or deliverables, in each case developed specifically for the Government referencing Government Data or Government design specifications under a contract or task/delivery order. Custom Development excludes any background intellectual property (e.g., the underlying supporting services, default configurations, and/or mechanisms), any generalized improvements, know-how, or techniques of general applicability, and anything developed independently without use of or reference to Government Data or Government design specifications.

Custom Model means a model created by fine-tuning, training, or otherwise adapting an LLM Developer-owned base model using Government Data or Custom Development. A Custom Model consists of the resulting model weights and parameters. It does not include the underlying base model, which remains the LLM Developer's background intellectual property. Custom Model excludes any background intellectual property (e.g., the underlying supporting services, default configurations, and/or mechanisms), any generalized improvements, know-how, or techniques of general applicability, and anything developed independently without use of or reference to Government Data or Government design specifications.

Data has the same meaning as defined in 44 U.S.C. §3502(16) and must, without limiting the generality of the foregoing, specifically include Data Inputs and Data Outputs.

Data Inputs means all Data, information, personally identifiable information (PII), or content submitted to the LLM and related operational systems by or on behalf of the Government, including but not limited to user prompts, queries, instructions, system prompts, source data, documents, knowledge bases, Government email addresses, user account information, and any other information or content submitted to the LLM and related operational systems by or on behalf of the Government. This definition specifically excludes any background intellectual property or information existing prior to entry into a contract or task/delivery order or developed independently

by such LLM Developer. Data Inputs excludes Background Data and technical system data (including logs, performance metrics, token counts, resource utilization and processing times). These exclusions only apply when the record does not capture the content of any Data Input or free-text or copied content which may contain Data Inputs.

Data Outputs means content generated by the LLM in response to Data Inputs in the performance of a contract, including but not limited to responses, results, analyses, anonymized data and synthetic or derivative content produced from Data Inputs. Data Outputs excludes Background Data and technical system data (including logs, performance metrics, token counts, resource utilization and processing times). These exclusions only apply when the record does not capture the content of any Data Input or Data Output or free-text or copied content which may contain Data Inputs or Data Output.

Government means any entity authorized to obtain procurement services through GSA pursuant to 40 U.S.C. § § 501, 502.

Government Data means Data Inputs and Data Outputs.

Government LLM System Administrator means the ordering agency government officials designated to oversee and manage that agency's provisioned tenant workspace or environment of the LLM.

Information system means a discrete set of information resources organized for the collection, processing, maintenance, use, sharing, dissemination, or disposition of information (44 U.S.C. 3502(8)).

Large Language Model Artificial Intelligence System (LLM) means a generative artificial intelligence model trained on vast, diverse datasets that enable the model to generate natural-language responses to user prompts (Ref. Executive Order 14319 Section 2). LLM includes the integrated technical and operational environment in which the model is configured, deployed, operated, monitored, or made available for the processing of Government Data, including the model, hosting and access infrastructure, system prompts, configurations, knowledge bases, retrieval mechanisms, Application Programming Interfaces (APIs), user interfaces, guardrails, monitoring systems, and associated workflows.

LLM Developer means the party that performs LLM Development Tasks (as defined in NIST AI RMF 1.0, Appendix A as "AI Development Tasks") by designing, developing, training, fine-tuning, calibrating, testing, publishing, licensing, or otherwise making available an LLM, including model weights, interfaces, model cards, safety documentation, or conditional usage restrictions. LLM Developer maps primarily to the "AI Development" actor category in NIST AI RMF 1.0, Appendix A.

LLM Service Provider means the party that performs LLM Deployment and Operation & Monitoring (as defined in NIST AI RMF 1.0, Appendix A as "AI Deployment and Operation & Monitoring") tasks by providing an LLM-enabled application, product, service, workflow, API, user interface, or business capability to Customers or End Users, and controlling the manner in which the service is presented, accessed, administered, supported, or used within that application or workflow. LLM Service Provider maps primarily to the "AI Deployment" and "Operation and Monitoring" actor

categories in NIST AI RMF 1.0, Appendix A.

LLM System Integrator means the party that performs LLM Design and LLM Deployment (as defined in NIST AI RMF 1.0, Appendix A as “AI Design” and “AI Deployment”, respectively) tasks by selecting, configuring, adapting, or materially controlling how an LLM System performs in a specific deployment or use case, including by selecting models, setting system prompts, prompt templates, Retrieval-Augmented Generation (RAG) sources, fine-tuning data, tools, plugins, agents, guardrails, filters, evaluation criteria, human-review thresholds, or output constraints. LLM System Integrator maps primarily to the "AI Design" and "AI Deployment" actor categories in NIST AI RMF 1.0, Appendix A, and includes parties NIST identifies as "system integrators."

LLM System Operator means the party that performs LLM Deployment and Operation & Monitoring (as defined in NIST AI RMF 1.0, Appendix A as “AI Deployment and Operation & Monitoring”) tasks by hosting, serving, operating, or providing access to an LLM or LLM System, including through cloud infrastructure, model endpoints, runtime environments, API availability, capacity management, logging, retention, and runtime security. LLM System Operator maps primarily to the "AI Deployment" and "Operation and Monitoring" actor categories in NIST AI RMF 1.0, Appendix A, and may also be a "Third-party entity" where it provides such services for another organization.

Material Change means any modification that could affect the trustworthiness, security, or operational integrity of the performance of the contract, especially as they pertain to the processing and protection of Government Data and the contractors involved in that process.

Personally Identifiable Information (*PII)* has the same meaning as defined in OMB Circular No. A-130.

(c) *Order of Precedence.* These terms and conditions establish specific requirements that take precedence over conflicting provisions in the LLM Developer policies, requirements, terms, conditions, or commercial agreements.

(d) *LLM Developer Responsibility for LLM*. The LLM Developer must:

(1) Exercise due diligence in overseeing the development of the model and performing the activities of an LLM Developer.

(2) Notify the Contracting Officer, within 72 hours, of any known non-adherence to these terms and conditions.

(3) The LLM Developer will:

(i) Include applicable requirements from these terms and conditions in agreements with LLM Developer(s), System Operator(s), System Integrator(s), and Service Provider(s) when performing services under government contracts; or

(ii) Obtain attestation when requested by the government, from an individual with appropriate authority representing LLM’s Developer(s), System Operator(s), System Integrator(s), and Service Provider(s) who are partners with the LLM Developers in performing services under government contracts, that requirements have been implemented and provide a new attestation if their compliance changes.

(4) Where meeting these requirements depends on how the LLM is configured or deployed by the Government, the LLM Developer will provide the guidance, documentation, and technical capability for the Government to implement and verify compliance upon request.

(e) *Intellectual Property Rights.*

(1) *Rights in Government Data.*

(i) The Government retains full ownership of, and will own, all Government Data and Custom Developments. The LLM Developer does not have any rights to use Government Data or Custom Developments other than those described in paragraph (e)(1)(iii).

(ii) Only LLM Developers, LLM System Operators, LLM System Integrators, and LLM Service Providers who are partners with the LLM Developers in performing services under government contracts may receive and process Government Data.

(iii) The LLM Developer is granted a limited, revocable, non-exclusive, non-transferable, worldwide, fully paid-up, royalty-free right and license to copy, store, transmit, modify, display, and use Government Data and Custom Developments for the duration of each individual awarded contract or task/delivery order solely for the following permitted purposes:

(A) Performing the specific requirements;

(B) Providing technical support and maintenance as required; and

(C) Providing such other uses as may be expressly authorized in writing by the Contracting Officer.

(iv) To the extent the LLM Developer obtains any intellectual property rights in Government Data, or any improvements, enhancements, or derivative works thereof, the LLM Developer assigns and transfers all such rights to the Government effective immediately upon creation. This assignment does not extend to the LLM Developer's background intellectual property, base models, or any generalized improvements, know-how, or techniques of general applicability, as excluded from the definition of Custom Development.

(v) The LLM Developer retains ownership of the underlying LLM, base models, and Background Data in its original form.

(2) *Right to Use to Government*. The LLM Developer grants to the Government an irrevocable, royalty-free, non-exclusive right to use the LLM for the duration of the work defined in the contract or task/delivery order. This right to use is strictly limited to the specific purposes and scope of work defined within the contract or task/delivery order and subject to the LLM Developer’s Terms of Service with this Addendum taking precedence in the event of a conflict, and is restricted to the commercially available features, Background Data, and functionality of the LLM as described in and necessary to fulfill a contract or task/delivery order and its accompanying documentation, unless specified otherwise. This right to use includes the right to:

(i) Operate and access the LLM through agreed-upon methods which includes LLM Developer’s acceptable use policy;

(ii) Allow authorized Government personnel and contractors to use the LLM; and

(iii) Integrate the LLM with Government systems.

(3) *Prohibited Uses of Government Data.* Use of Government data is limited per paragraph (e)(1)(iii). Other uses of Government data are prohibited. Examples of prohibited use of Government Data includes:

(i) Training or fine-tuning an LLM, including those operated by third parties, or to develop the LLM(s) for any other customers or any commercial or non-commercial purposes unless LLM Developer and Government contract for a Custom Development.

(ii) Using Government Data to inform the LLM Developer's advertising, marketing, sales, monetization, strategy, operations, or other business decisions, or to provide to other Government or non-Government entities.

(iii) Retaining, accessing, or using beyond the scope and duration expressly permitted in the contract including LLM Developer's standard data retention policies.

(iv) Processing or storing Government Data with, or transferring Government Data to, any party not authorized under a contract or task/delivery order or without appropriate extension (flowdown) of applicable requirements.

(v) Selling or licensing Government Data to any party;

(vi) Knowingly reconstructing Government Data from technical system data (as defined in Data Inputs and Data Outputs).

(4) *Government Data Handling and Processing Requirements*. Unless otherwise expressly authorized by the Contracting Officer:

(i) The LLM Developer must implement reasonable technical, administrative, physical, and organizational safeguards to protect Government Data from loss, damage, destruction, unauthorized alteration, or corruption, and prevent its unauthorized, accidental, or unlawful access, disclosure, use, or processing;

(ii) The LLM Developer must implement automated processing systems and operational controls to minimize human review of and access to Government Data and make such review and access visible to the Government. For every product that processes Government Data, whether or not content is retained, the LLM Developer must:

(A) permit human access to Government Data only to provide the service, or respond to an operational incident.

(B) require approval by a second person on a written justification before any access to production systems, and record each such grant, including the requester and the reason, in an audit log; and

(C) bar all other staff access to Government Data except as set forth in (A) above.

(iii) For each product used to process Government Data, the LLM Developer must provide to the Contracting Officer the highest level of protection in (A) through (D) below that the product supports, and must state per product which level it provides.

(A) Where an LLM Developer's product supports zero data retention, the LLM Developer must offer it and, where the Government elects it, must not retain Government Data at

rest after the response is returned. The parties acknowledge that a single API workspace cannot operate under both zero data retention and a Government-held encryption key; where content must be stored, the Government may elect a separate workspace under its own key.

(B) Where an LLM Developer's product supports access reporting (Access Transparency), the LLM Developer must enable it for the Government and report human views of retained Government Data through the Compliance API within two business days, identifying the team and the reason.

(C) Where a product does not support access reporting, the LLM Developer must, as a compensating control, restrict human access to Government Data to a defined list of named, authorized personnel reaching it through a controlled mechanism, provide that list and updates to the Government, and identify the individuals who accessed specific Government Data upon Government request. At a minimum, the compensating control will include a process whereby human query or review of Government Data must be logged and made available to the Government on request.

(D) Where the LLM Developer can offer neither access reporting nor the compensating control for a product, it must use best efforts to log human review and access and to respond to Government requests for that information, and must provide a written plan, on a timeline it proposes for implementing (A) through (C) in this section. Best efforts under this paragraph are not sufficient for any product that retains Government Data thirty (30) days beyond the Government's control or the term of the contract, task/delivery order, or BPA; such a product must provide zero data retention under (A) or the compensating control under (C).

(iv) For each product used to process Government Data, the LLM Developer must include in its documentation under paragraph (f)(1) whether human access to Government Data is limited to persons located in the US or subject to its laws.

(v) Human review or access permitted under this section does not include any right to export, copy out, or transfer Government Data. The LLM Developer must maintain controls so that Government Data may be viewed for a permitted purpose but not exported or transferred outside the service environment.

(vi) The LLM Developer must notify the Contracting Officer and/or Government LLM System Administrator in advance of any change that lowers a product's level of protection under this section, including any change to zero data retention, access reporting, or encryption-key settings. Improvements to protections under this section may be made at any time.

(vii) The LLM Developer must ensure Government Data is stored or processed only when reasonably necessary for the performance of the contract;

(viii) The LLM Developer must provide tools that enable the Government to maintain detailed records of all processing activities involving Government Data;

(ix) The LLM Developer must comply with the following Data localization requirements, including

but not limited to:

(A) Not removing or allowing removal of any such Government Data from the agreed-upon premises or FedRAMP-authorized services, without express written consent from the Contracting Officer or the Government LLM System Administrator; and

(B) Not transmitting, storing, taking, or accessing such Government Data, or allowing such Government Data to be transmitted, stored, taken, or accessed by any means outside of the agreed-upon premises or authorized services, without express written consent from the Contracting Officer or an affirmative action taken by the Government LLM System Administrator. For clarity, the requirements in this subsection (v) are not applicable where the Government Agency LLM Administrator has the ability to select Data localization options such as FedRAMP Certified/Authorized, US-only, vs Global. The Government Agency LLM Administrator is the responsible party where such features and options are available.

(x) The LLM Developer must implement and maintain appropriate technical and organizational measures to ensure that all such Government Data is logically segregated from the Data of any non-Government customer or client, and is not commingled with Data of other customers or clients and provides adequate defense in depth through access controls, policy enforcement points, labeling, and/or encryption mechanisms and must perform continuous monitoring in accordance with the products certifications to protect against unauthorized access by external and internal threat actors, third parties, or disclosure resulting from human or machine error. Logical segregation does not require physical or dedicated instances within contractor-hosted or controlled systems; continuing compliance with the specified FedRAMP authorization level is deemed to satisfy these requirements by achieving equivalent security and separation.

(xi) Upon completion, termination or expiration of the contract or task/delivery order, unless otherwise directed in writing by the Contracting Officer, the LLM Developer must securely and permanently delete all such Government Data and any Custom Developments from the LLM and all its other systems and all copies, backups and derivatives thereof, and certify deletion to the Contracting Officer in writing. LLM Developer may retain Government data residing in LLM Developer's routine backup or archival media which shall be overwritten or destroyed in accordance with LLM Developer's standard backup retention and media sanitization cycle. Notwithstanding the foregoing, LLM Developer may retain Government Data to the extent and for the duration required by applicable U.S. law, regulation, court order, litigation hold, or provided that any Government data (a) remains subject to the all requirements of this Addendum, (b) used solely for the purpose giving rise to the retention obligation as allowed by this Addendum, and (c) promptly deleted upon the expiration of that obligation.

(xii) *Custom Developments and Custom Model Rights.* If there is a requirement for Custom Developments or Custom Models as deliverables in the contract, the LLM Developer must:

(A) Dedicate the Custom Developments or Custom Models, to the Government's exclusive use;

(B) Treat such Custom Developments, Custom Models, and all associated Data as the Government's confidential information; and

(C) Not use, reproduce, or derive benefit from such Custom Developments or Custom Models for any other purpose, or for the benefit of any other party, without express written authorization from the Contracting Officer.

(xiii) *Feedback.* LLM Developer may use feedback provided through a general customer service

channel such as but not exclusively email or a support ticket to improve its systems. If feedback includes marked Government sensitive or confidential information, that feedback may not be used for system improvement purposes or any other purposes (except solely in the performance of a contract). LLM Developer will not be responsible for maintaining the confidentiality of unmarked Government sensitive or confidential information when such information has been provided through a general customer service channel such as but not exclusively, email or support ticket.

(f) *LLM Developer Obligations for Compliance, Reporting, and Documentation.* The LLM Developer must:

(1) Disclose all LLMs used or made available under government contracts. Disclosures are due to the Contracting Officer or ordering Contracting Officer by the date specified in the contract or task/delivery order. If no date is specified, disclosures are due within 120 days after commencing work under the contract or task/delivery order.

(2) Maximize the use of LLMs that meet the following criteria:

(i) Controlling Entity & Jurisdiction: Each LLM is developed, managed, and operated by an entity that is incorporated in the United States (U.S.) and is subject to U.S. law and jurisdiction.

(ii) Protection Against Foreign Compulsion:

(A) Each LLM's performance, operations, or protections cannot be controlled by a foreign Government.

(B) No foreign Government or entity can compel disclosure of Government Data, or operational control that could compromise each LLM's integrity or security.

(C) Each LLM, and any components performing core model, data storage or processing, output generation, or security functions, are prohibited from being developed, managed, or operated by entities subject to the direction, influence, or control of adversary foreign governments (see 15 CFR 791.4).

(iii) Component Flexibility & Risk Mitigation: Incidental foreign-developed components (e.g., open-source components, published research), ancillary Third-party services, or globally operated infrastructure dependencies are permissible, provided:

(A) They do not introduce security risks or foreign control that would violate criteria (f)(2)(i) or (f)(2)(ii).

(B) The systems storing or processing Government Data satisfy applicable Federal security requirements, including those for data residency and isolation.

(C) A risk-based approach is applied, focusing on objective criteria such as ownership, control, hosting, and security posture.

(3) *Intermediary Processing.*

(i) For Claude Enterprise and Claude for Government, if the LLM uses intermediary processing such as reasoning, retrieval, or agentic processes, the LLM Developer will commit to implement and/or develop tools which provides the Government a means to implement human oversight, intervention, and traceability through audit trail and one or a combination of the following: customer-configured logs or the user interface. At minimum, the LLM must include:

(A) Summarized intermediate processing actions and decision points from input to output, to the extent the LLM performs and can record them;

(B) The model generating the response;

(C) Data retrieval methods employed (e.g., Retrieval-Augmented Generation (RAG), web search, tool calls, agent actions), with source attribution with direct links, where available, to allow a consuming user or application to identify materials used in generating the response.. For clarity, non-substantive context signals, such as approximate location, date, time, locale, or weather, need not be attributed.

(ii) For first party and third party API, if the LLM uses intermediary processing such as reasoning, retrieval, or agentic processes, the LLM Developer will commit to implement and/or develop tools which provides the Government a means to implement human oversight, intervention, and traceability through API response. At minimum, the LLM must include:

(A) Summarized intermediate processing actions and decision points from input to output, to the extent the LLM performs and can record them;

(B) The model generating the response;

(C) Data retrieval methods employed (e.g., Retrieval-Augmented Generation (RAG), web search, tool calls, agent actions), with source attribution with direct links, where available, to allow a consuming user or application to identify materials used in generating the response. For clarity, non-substantive context signals, such as approximate location, date, time, locale, or weather, need not be attributed.

(4) Notify the Contracting Officer and a designated Government security point of contact as soon as possible but not longer than within 72 hours of the discovery of any incident (as defined by the Federal Information Security Modernization Act of 2014 in 44 U.S.C. § 3552(b)(2)) affecting any contractors, including third-party entities, handling Government Data and provide daily status updates to all points of contact until resolved.

(i) Such notification must include, to the extent known at the time:

(A) Nature and scope of the incident;

(B) Data potentially affected;

(C) Immediate remediation steps taken;

(D) Timeline for full resolution; and

(E) Measures to prevent recurrence.

(ii) Where FedRAMP incident communication and response procedures conflict with contractual requirements, the FedRAMP procedures take priority if the system is required to be FedRAMP Authorized.

(iii) The LLM Developer must preserve all relevant logs, forensic images, and incident artifacts for a minimum of 90 calendar days from a security incident involving Government Data to support follow-on investigation activity by law enforcement entities.

(5) Provide, upon Government request, and under appropriate confidentiality protections, existing commercial documentation or disclosures sufficient to demonstrate compliance, including but not limited to:

(i) Verification of compliance with these terms and conditions and the contract;

(ii) List of entities involved in performance of a contract or task/delivery order, the entities' role(s), which requirements of these terms and conditions are applicable to each entity, and the entities' approach to addressing the requirements in these terms and conditions;

(iii) Known LLM decision-making processes, and operational parameters that materially affect outputs, as documented or as configured.The contractor is not required to disclose internal confidential information, proprietary source code, model weights, or trade secrets.

(iv) LLM Developers, general model characteristics, intended use, limitations, and risk considerations to enable informed and responsible use by the Government. The contractor is not required to disclose proprietary source code, model weights, or trade secrets.

(v) System documentation consistent with NIST AI Risk Management Framework guidelines, Unbiased AI Principles, and LLM Transparency requirements such as system cards or equivalent documentation.

(vi) The LLM Developer must provide commercially available means for the Government to implement appropriate human oversight, intervention, and traceability for the contracted use case, and to understand the role of various components in generating responses, to the extent commercially available and technically feasible;

(vii) Privacy controls effectiveness and PII processing prohibition capabilities;

(viii) Testing methodologies used to detect and mitigate noncompliance with the unbiased AI principles described in paragraph (j)(1);

(ix) Known biases (including commercial, political, or personal considerations, advertising, endorsements, or fraudulent/corrupt interests), limitations, truthfulness concerns, and performance metrics;

(x) Influence, direction, or control of an adversary foreign governments (see 15 CFR 791.4);

(xi) FedRAMP authorization packages as appropriate, continuous monitoring data, and assessment artifacts necessary to issue or maintain an authorization to operate, consistent with 44 U.S.C. § 3609(a)(8) and OMB M-24-15 and successor requirements; and

(xii) Any other information necessary for the Government to monitor and evaluate the LLM's performance, risks, and effectiveness and to complete an LLM Impact Assessment required by OMB M-25-21 and successor requirements.

(g) *Privacy and Confidentiality Protections.* To the extent the LLM Developer has available tools, they must provide tools to enable the Government to implement appropriate Government-configurable controls, including but not limited to automated detection mechanisms and clear user notifications to manage the entry or persistence of PII within the LLM.

(h) *Data Portability and Interoperability.* The LLM Developer must:

(1) Provide and ensure that LLM services support the export of Government Data (including but not limited to: user-generated content, interaction history, uploaded content and media, and knowledge bases) in open or standard machine-readable formats, together with sufficient associated metadata and schema information necessary for the Government to retrieve, use, or migrate Government

Data to another service.

(2) Provide and support documented APIs sufficient for Government integration and data export to the extent the service uses APIs.

(3) Not impose proprietary formats, technical restrictions, additional costs, or additional licensing conditions that materially impair the Government's ability to retrieve, use, or migrate Government Data to another service.

(4) Provide tooling or interfaces to facilitate migration of Government Data and Custom Developments in a secure manner.

(i) *Change Notification.*

(1) *Advance Notice of Specified Changes.* The LLM Developer must provide written notice to the designated Contracting Officer at least thirty (30) calendar days prior to:

(i) Adding, replacing, or materially changing any LLM System Operator(s);

(ii) Changes to the services handling Government Data, including discontinuation or material reduction of data protection controls applied to Government Data, or change in FedRAMP Authorization Status; or

(iii) Any modification, configuration, or training change to the LLM made to comply with any foreign government statute, regulation, or policy where that change materially affects model outputs, accuracy, or truthfulness.

(iv) This does not include regular model version changes as outlined in i(2).

(2) *Model Version Changes.* The LLM Developer must use reasonable efforts to provide the Government with concurrent access to any successor LLM, and associated documentation, for a minimum evaluation period of thirty (30) calendar days for major versions, and fifteen (15) calendar days for minor versions, prior to discontinuing or replacing the immediate prior LLM in use under a contract or task order.

(3) *Safety and Performance Degradation.* The LLM Developer must notify the Contracting Officer within ten (10) calendar days of identifying any material change to the Government deployed model’s then current Safeguards or configuration, where Safeguards means the safety classifiers, safety-related system prompt instructions, and usage policy enforcement controls that the LLM Developer applies to the Government’s deployment, that materially increases output bias, materially reduces the effectiveness of Safeguards, or materially degrades performance or truthfulness of outputs. The notification shall describe the change, its purpose, and where reasonably feasible and consistent with LLM Developer’s legal obligations the evaluation approach used and any new limitations, trade-offs, or potential impacts on outputs identified. For the avoidance of doubt, the paragraph does not apply to (i) changes intended to reduce unwarranted refusals or false positives or (ii) settings selected by or configurable by the Government LLM System Administrator. The definition of Material Change in paragraph (b) does not apply to this paragraph.

(4) *Emergency and Unplanned Changes.* If the LLM Developer implements a change without providing sufficient advance notice, or if a change is necessitated by a security or service

emergency, the LLM Developer must:

(i) Notify the Contracting Officer as soon as practicable upon learning of the change; and

(ii) Describe any remediation or rollback actions available.

(5) *Notification*. All notifications under paragraph (i) must include, to the extent known at the time of notification:

(i) Description of change;

(ii) Affected roles, services, systems, and Government Data;

(iii) Any limitations, trade-offs, or negative impacts;

(iv) Evaluation approach used to assess the change;

(v) Any remediations or rollback actions; and

(vi) Confirmation of applicable requirements of these terms and conditions have been flowed down appropriately.

(vii) Provide general information sufficient for the Government to take appropriate action, and any sensitive or confidential information to authorized Federal personnel. The Government, to the extent permitted by law and regulation, will safeguard and treat information obtained from the LLM Developer as confidential where the information has been marked "confidential" or "proprietary" by the LLM Developer. To the extent permitted by law and regulation, such information will not be released by the Government to the public pursuant to a Freedom of Information Act request, 5 U.S.C. § 552, without prior notification to the LLM Developer and opportunity to object to disclosure. The Government may transfer documents provided by the LLM Developer to any department or agency within the Executive Branch if the information relates to matters within the organization’s jurisdiction.

(j) *Performance, Evaluation, and Remediation.*

(1) *Unbiased AI principles*. The LLM Developer must ensure the LLM is developed and monitored in accordance with the following unbiased AI principles:

(i) The LLM Developer must use commercially reasonable efforts to design and train the LLM to respond to user prompts seeking factual information or analysis, prioritize accuracy, scientific inquiry, and objectivity and must acknowledge uncertainty where reliable information is incomplete or contradictory.

(2) *Government Evaluation Rights and Remediation*.

(i) The Government reserves the right to conduct automated assessments of the LLM, as deployed and configured for government users, at any time using its own benchmarks. These evaluations may assess bias, truthfulness, safety, unsolicited ideological content, and other factors determined by the Government in order to facilitate evaluations.

(ii) The LLM Developer must provide tools and interfaces that enable the Government to run its benchmarks in an automated fashion to test the production LLM and identify any material gaps.

(iii) The Government may provide such results to the LLM Developer to support remediation efforts.

(iv) All benchmarks, test data, and methodologies developed or used by the Government for such assessments are considered Government Data. The Government is under no obligation to disclose or provide access to the underlying data, methodologies, or systems, with the exception of data or methodologies used as the basis for an adverse action under paragraph (j)(3)(i).

(3) *Non-Compliance*.

(i) The Government retains the right to suspend use of the LLM until performance issues are satisfactorily addressed.The Government must disclose information, under appropriate confidentiality protections, to enable the LLM Developer to understand the basis for the Government’s determination and take reasonable remediation actions.

(ii) The LLM Developer retains the right to suspend the Government Agency's access, or individual users, in response to suspected violations of agreements or terms of use not outlined here and the Government accepts that LLM Developer will take actions to suspend based solely on automated trust and safety controls. The LLM Developer must provide best efforts to notify the Government prior to suspension and disclose related information to trust and safety scores, under appropriate confidentiality protections, to enable the Government to understand the basis for the LLM Developer's determination and to enable the Government to take reasonable remediation actions to restore access within 72 hours except when LLM Developer reasonably believes there is a risk or imminent attack to the LLM Developer’s services or the LLM Developer has documented findings of the Government Agency systematically violating the LLM Developer’s policies as set forth in the terms of use. The Government will work with the LLM Developer to provide relevant information and sample data or prompts when there is a dispute over the reason for suspension. In the event of repeat suspensions or an inability to remediate, the LLM developer may require the Government to disclose additional usage data in order to restore services. Should the Government fail to provide the information set forth above for repeated suspensions, LLM Developer is not required to restore access to the suspended user(s) and/or workspace.