Customer Contracting Toolkit for B2B AI Services

A Model Agreement Framework for Operating Companies Procuring AI Products and Services[1]

Intellectual Property Owners Association (IPO)

AI, Data & Emerging Technologies Committee

DRAFTER'S NOTE

About This Toolkit. The Customer Contracting Toolkit for B2B AI Services provides model contract language for operating companies of any type contracting for AI products or services (focused on software products or SaaS/hosted services) from a vendor. It is intended for use by in-house IP counsel as a starting point for negotiating AI procurement agreements. Provisions are drafted from the Customer's perspective, with fallback positions and drafter's notes throughout.

Parties. Throughout this Toolkit, the procuring operating company is referred to as "Customer" and the AI product or service vendor is referred to as "Provider."

Disclaimer. This Toolkit is provided for educational and informational purposes only. It does not constitute legal advice and does not create an attorney-client relationship.

[1] Co-authors:
Alex Bridge (Hewlett Packard Enterprise),
Shruti Costales (Shruti Law PLLC®),
Elias Domingo (Xylem),
Sam Gee (Kenvue),
Steven Hanley (Caterpillar Inc,),
Anne Marie Longobucco (Fenwick & West LLP),
Justin Mullen (Capital One),
Laura Nguyen (Hewlett Packard Enterprise), and
Kirk Sigmon (KellDann Law PLLC).

Table of Contents

About This Toolkit

This Toolkit provides model contract language for the procurement of AI products and services in a B2B context. It is organized into a Definitions section, thirteen substantive sections addressing distinct legal topics arising in AI transactions, and an exhibit template for AI-specific service level agreements. Each section is drafted from the Customer's perspective and includes operative provisions, drafter's notes, and fallback positions where applicable.

The Toolkit is intended for in-house counsel. It assumes familiarity with IP concepts including license grants, work-for-hire doctrine, field-of-use restrictions, indemnification, and trade secret law.

How to Use This Toolkit

Opening Positions: Black text represents Customer-favorable opening positions for negotiation.

Drafter's Notes: Blue-bordered boxes provide legal context, analysis, and guidance on key issues and risks.

Alternative/Fallback Positions: Amber-bordered boxes provide more balanced language a sophisticated Provider might accept, each with identified risks.

Placeholders: Gray-bordered boxes mark sections that remain to be drafted.

Bracketed Items: Items in [brackets] require customization to the specific transaction.

Definitions

The following defined terms apply throughout this Toolkit. Additional defined terms appear within individual sections; those terms are defined where they first appear and are incorporated herein by reference.

"Agreement"	The agreement between Customer and Provider into which the provisions of this Toolkit are incorporated, together with all Exhibits, Schedules, Order Forms, and Statements of Work incorporated therein.
"AI Model" or "Model"	The artificial intelligence or machine learning model, including all associated model weights, parameters, embeddings, architectures, and inference logic, made available by Provider to Customer under this Agreement, whether accessed via API, hosted interface, embedded software, or other delivery mechanism.
"AI Security Incident"	Any AI-specific security incident affecting the AI Services or AI Model, including without limitation: (a) prompt injection attacks; (b) model inversion, extraction, or stealing attempts; (c) unauthorized access to model weights, embeddings, or fine-tuning data; (d) cross-customer data contamination through shared model infrastructure; (e) unintended memorization or reproduction of Personal Data or Customer Data in model Outputs; or (f) other attacks targeting AI-specific attack surfaces. For the avoidance of doubt, an AI Security Incident may also constitute a Security Event.
"AI Services"	The AI-powered products, software-as-a-service offerings, APIs, interfaces, and related services provided by Provider to Customer under this Agreement, as described in the applicable Order Form or Statement of Work. For the avoidance of doubt, AI Services encompasses the full technical stack through which the AI Model is made available to Customer, including the AI Model, APIs, user interfaces, pre- and post-processing components, and associated infrastructure.
"Authorized User"	Any employee, contractor, or agent of Customer who is authorized by Customer to access and use the AI Services under this Agreement.
"Base Foundation Model"	Any general-purpose AI or machine learning model developed or licensed by a third party (including without limitation large language models, multimodal models, or other foundation models) upon which Provider's AI Model (as defined above) is built, layered, or otherwise dependent, in whole or in part.

“Confidential Information”	All non-public information disclosed by or on behalf of one party to the other, whether oral, written, or electronic, that is designated as confidential or that a reasonable person would understand to be confidential given the nature of the information and circumstances of disclosure. For the avoidance of doubt, Customer’s Confidential Information includes all Customer Data, Inputs, and Customer-Specific Customizations, regardless of designation.
"Customer"	The operating company identified as the customer or subscriber in the applicable Order Form or Agreement cover page.
"Customer Data"	All data, content, and information that Customer owns or controls and that is provided to Provider or to the AI Services by or on behalf of Customer in connection with this Agreement, including without limitation Inputs, data used for Customer-Specific Fine-Tuning, and configuration data. Customer Data is the broad umbrella term; Inputs is the narrower operational term for data submitted at inference time.
"Customer-Derived Improvements"	Any modification, enhancement, optimization, capability, or performance improvement to any AI Model that was enabled by, trained on, derived from, or informed by: (a) Customer Data; (b) Customer-Specific Customization; (c) usage patterns or interactions arising from Customer’s use of the AI Services; or (d) any combination of the foregoing. For the avoidance of doubt, Customer-Derived Improvements does not include improvements derived solely from Feedback.
“Customer Materials”	Any text, photos, graphics, recordings, documentation, images, drawings, designs, plans, artwork, models, specifications, processes, software, algorithms, or other work product (including all related drafts and preparatory materials) created by Customer using the AI Services, together with Customer's Confidential Information and proprietary know-how, excluding any Outputs or Deliverables allocated under this Agreement. Customer Materials are owned exclusively by Customer and Provider acquires no title to or ownership of any Customer Materials.
"Customer-Specific Configuration"	Any system instructions, system prompts, parameter settings, hyperparameter adjustments, or other configuration customizations applied to the AI Model that are specific to Customer and not made available to other customers of Provider.
"Customer-Specific Customization"	Collectively, Customer-Specific Configuration and Customer-Specific Fine-Tuning.
"Customer-Specific Fine-Tuning"	Any fine-tuning, retraining, reinforcement learning from human feedback (RLHF), or other adaptation of the AI Model’s weights or parameters performed using Customer Data or specifically for Customer.

"Deliverables"	Outputs that are specifically developed for Customer pursuant to a Statement of Work or Order Form, as distinct from generic Outputs of the AI Model.
“Effective Date”	The date on which this Agreement becomes effective, as specified on the Agreement cover page or in the applicable Order Form. Unless otherwise specified, the Effective Date is the date of last signature by both parties.
"Feedback"	Suggestions, ideas, recommendations, comments, or other information provided by Customer or its Authorized Users to Provider relating to the AI Services, AI Model, or Documentation. For the avoidance of doubt, Feedback does not include Customer Data, Inputs, or Customer-Derived Improvements.
"Inputs"	The data, content, prompts, queries, and other information submitted by Customer or its Authorized Users to the AI Services at inference time for the purpose of generating Outputs.
"Losses"	Any and all claims, demands, suits, proceedings, damages, fines, penalties, costs, and expenses (including reasonable attorneys’ fees and litigation costs).
"Model Documentation"	All technical documentation, model cards, system cards, API specifications, configuration guides, and other materials provided by Provider describing the AI Model’s architecture, capabilities, limitations, intended use, and known risks.
"Model Intellectual Property"	All intellectual property rights (including patents, copyrights, trade secrets, and know-how) subsisting in or relating to: (a) the AI Model; (b) the Provider’s Proprietary Model; (c) model weights, parameters, and architectures; and (d) training methodologies and processes used to develop any of the foregoing, to the extent owned or licensable by Provider.
"Outputs"	Content, data, predictions, recommendations, classifications, or other results generated by the AI Model in response to Inputs.
"Personal Data"	Any information that identifies, relates to, describes, is reasonably capable of being associated with, or could reasonably be linked, directly or indirectly, to an identified or identifiable natural person, including as defined under applicable privacy and data protection law. For the avoidance of doubt, Personal Data includes "personal information" as defined under the California Consumer Privacy Act (as amended by the California Privacy Rights Act) and "personal data" as defined under the General Data Protection Regulation (EU) 2016/679 and UK GDPR.
"Provider"	The AI product or service vendor identified as the provider or supplier in the applicable Order Form or Agreement cover page.
"Provider’s Proprietary Model"	Any version of an AI Model that Provider has independently fine-tuned, retrained, adapted, or otherwise modified from a Base Foundation Model using Provider’s own data and resources,

	exclusive of any Customer Data or Customer-Specific Customization, and which Provider makes generally available to its customers.
"Security Event"	Any actual or reasonably suspected unauthorized access to, acquisition of, use of, disclosure of, modification of, or destruction of Customer Data, or any actual or reasonably suspected breach of the security of any Provider system that stores, processes, or transmits Customer Data, consistent with applicable breach notification law.
"Subprocessor"	Any third party engaged by Provider that accesses, processes, stores, or transmits Customer Data in connection with the provision of the AI Services, including without limitation, cloud infrastructure providers, model providers, and downstream service providers.
“Term”	The period during which this Agreement is in effect, beginning on the Effective Date and continuing until the expiration or earlier termination of this Agreement in accordance with its terms, as specified in the applicable Order Form or Agreement cover page.
"Third-Party Intellectual Property Rights"	Any intellectual property rights (including patents, copyrights, trademarks, trade dress, trade secrets, and rights in confidential information and know-how) held by any party other than Customer or Provider.
"Training Data"	The data, content, and information used by Provider to train, develop, or improve the AI Model or Provider’s Proprietary Model, excluding Customer Data and Customer-Specific Fine-Tuning data. For the avoidance of doubt, Training Data is Provider’s responsibility and Customer has no visibility into or control over its composition except as expressly disclosed pursuant to Section 4.1.5.

DRAFTER'S NOTE

On Customer Data vs. Inputs. These two terms are intentionally distinct. “Customer Data” is the broad umbrella covering all Customer-originated data flowing through the AI Services, including data used for fine-tuning, configuration, and inference. “Inputs” is the narrower operational term for data submitted at inference time to generate Outputs. The distinction matters most in Section 6 (use restrictions on Outputs, which attach to inference-time activity) and in Section 4 (where fine-tuning data and inference inputs are subject to different restrictions). Section 7 addresses liability allocation more broadly, including for Training Data-related claims addressed separately in Section 8.

On Outputs vs. Deliverables. All Deliverables are Outputs, but not all Outputs are Deliverables. Deliverables are Outputs created specifically for Customer pursuant to a particular engagement (typically a Statement of Work), as distinct from generic model results that any customer might receive. The distinction matters for ownership: under Sections 1–3 and Section 5, Deliverables are owned exclusively by Customer, while broader Output ownership is addressed under different terms.

On Feedback vs. Customer-Derived Improvements. These terms address distinct categories of Customer contribution: Feedback is suggestions and ideas Customer communicates to Provider; Customer-Derived Improvements are technical improvements to the model derived from Customer Data, configurations, or usage. They have different ownership treatment in Sections 4 and 9.

On Security Event vs. AI Security Incident. Security Event covers traditional data breaches (unauthorized access, disclosure, modification, destruction of Customer Data) and triggers traditional breach notification obligations. AI Security Incident covers AI-specific attack vectors (prompt injection, model inversion, embedding leakage, cross-customer contamination) that may not trigger traditional breach notification but warrant separate response obligations. The two categories may overlap; they are not mutually exclusive.

Section 1: Ownership of Customer's Data Input into AI

This Section addresses ownership of Customer Data submitted to the AI Services. The opening position reserves all ownership and IP rights in Customer Data to Customer. The provisions below apply regardless of the delivery mode of the AI Services (SaaS, API, or on-premises) and operate alongside the model license provisions in Section 4.

Section 1.1: Customer Ownership of Inputs and Customer Data

All Customer Data, including all Inputs submitted by or on behalf of Customer to the AI Services, shall remain the sole and exclusive property of Customer. As between the parties, Customer retains all right, title, and interest, including all intellectual property rights, in and to the Customer Data. Provider acquires no ownership rights or other proprietary interests in the Customer Data by virtue of this Agreement or its provision of the AI Services.

Provider shall treat all Customer Data as Customer's Confidential Information under this Agreement and shall not use, reproduce, disclose, or distribute Customer Data except as expressly permitted by this Agreement or as otherwise authorized in writing by Customer.

DRAFTER'S NOTE

Customer Ownership Is the Opening Position. In sophisticated B2B AI procurement, Customer ownership of input data is the near-universal opening position and is almost always accepted by Provider. The substantive negotiation tends to occur not on ownership itself, but on (a) the scope of permitted Provider use (Section 2 below); and (b) carve-outs for aggregated or anonymized data (also Section 2). Counsel should not over-negotiate ownership at the expense of leverage on use restrictions, where the real exposure lies.

Confidentiality Treatment. Treating Customer Data as Confidential Information is important because ownership alone does not regulate Provider's permitted handling. The confidentiality framework typically imposes specific safeguarding obligations, restrictions on access, and survival provisions that ownership alone does not. Counsel should confirm the Confidentiality provisions of the broader Agreement adequately address Customer Data.

Section 2: Use of Customer's Input Data; Choice of Data

This Section governs Provider's use of Customer Data. It restricts use of Customer Data for training, fine-tuning, or improving any AI Model that benefits any party other than Customer, and provides limited carve-outs for permitted uses. This Section operates in conjunction with Section 4.3 (Restrictions on Use of Customer Data in Model Development), which addresses similar restrictions from the model license perspective.

Section 2.1: Prohibition on Use of Customer Data for AI Training

Provider shall not use any Customer Data — including any Inputs, Deliverables, Personal Data, or other information submitted by or on behalf of Customer, and including any such information that has been de-identified, aggregated, or anonymized — for the purpose of training, fine-tuning, retraining, validating, benchmarking, developing, evaluating, or distributing any AI Model, system, or service of Provider or any third party, except as expressly permitted under Section 2.2.

DRAFTER'S NOTE

Why the De-Identified / Aggregated / Anonymized Carve-In Matters. Vendors routinely seek to use "de-identified," "aggregated," or "anonymized" data for training on the theory that such data no longer raises privacy concerns. The opening position above blocks this carve-out because: (a) re-identification of supposedly anonymized data is increasingly feasible; (b) even truly anonymized Customer Data may reveal competitive information, usage patterns, or trade secrets; and (c) once data enters a training pipeline, it may be irrecoverable from model weights.

Higher-Risk Data Categories. This opening position is particularly important where Customer Data includes: (i) Personal Data; (ii) trade secrets or confidential business information; (iii) information subject to regulatory restrictions (HIPAA, GLBA, FERPA, etc.); or (iv) information whose competitive value depends on exclusivity. For lower-risk categories of data, Customer may have flexibility to accept the more permissive language in the fallback below, but the analysis should be conducted on a data-category-by-data-category basis and in accordance with the risk profile of the organization.

FALLBACK — LIMITED PERMISSION FOR ANONYMIZED DATA

Where Customer's data risk profile permits more flexibility, the following fallback may be acceptable:

"Notwithstanding Section 2.1, Provider may use Customer Data that has been irreversibly anonymized (such that re-identification is technically infeasible using reasonable means) for the limited purpose of improving the AI Services generally; provided that (a) such use shall not include any Personal Data, Customer Confidential Information, or Customer trade secrets; (b) Provider shall not attempt to re-identify Customer or any individual from such data or perform any action that would allow a third party to re-identify Customer or any

individual from such data; and (c) such use shall not confer any competitive advantage on any other specific customer of Provider derived from Customer's use of the AI Services."

Risk: Anonymization is rarely truly irreversible; sophisticated re-identification techniques continue to advance. "Reasonable means" is doing significant work in this provision and should be paired with a recognized anonymization standard (e.g., NIST SP 800-188) where possible. Counsel should also consider pairing this fallback with audit rights under Section 11 and indemnification under the Agreement.

Section 2.2: Permitted Internal Use

Notwithstanding Section 2.1, Provider may use Customer Data solely as necessary to:

- provide, maintain, support, and operate the AI Services for Customer pursuant to this Agreement;
- perform Customer-Specific Fine-Tuning expressly requested by Customer in writing, subject to the ownership and use restrictions in Section 4, where the result of the Customer-Specific Fine-Tuning is available is limited to Customer and not available to third-parties;
- comply with applicable law, regulation, or valid legal process, provided that Provider provides Customer with prompt prior notice of any such legal requirement to the extent permitted by law; and
- respond to security incidents, prevent abuse, or address operational issues directly affecting Customer's use of the AI Services.

DRAFTER'S NOTE

Note on Scope. These permitted uses are intentionally narrow and not all will be appropriate for each use case. Counsel should review the list in light of Customer's specific use case — e.g., if Customer expects Provider to perform analytics or generate reports on Customer's usage, those use cases should be added to this list or addressed in a separate Statement of Work.

Internal AI Use by Provider. Provider may push to use Customer Data in "secure internal AI applications" that purportedly do not affect any other customer. Counsel should be skeptical of this position: "internal" AI use can still create training data exposure, retention concerns, and audit complications. If Customer accepts internal AI use, the permission should be tightly scoped, with confidentiality obligations explicitly extended to such use and an express prohibition on incorporating such data into any external model.

Section 2.3: Customer Control over Data Selection

Customer shall have sole discretion to determine which Customer Data, if any, is provided to the AI Services. Provider shall not require Customer to provide any particular type or category of Customer Data as a condition of accessing or using the AI Services, except to the extent strictly necessary for the operation of the AI Services as described in the Order Form or Statement of Work. Customer may at any time withhold, modify, or delete Customer Data, and Provider shall promptly delete or return such Customer Data in accordance with Customer's written instructions, subject only to retention obligations expressly required by applicable law.

DRAFTER'S NOTE

Why This Section Matters. Vendors increasingly design AI Services to capture broad categories of Customer Data — often by default — on the theory that more data improves model performance. This Section places control of data selection with Customer, consistent with privacy principles of data minimization and purpose limitation. Counsel should also confirm whether Provider provides granular controls (e.g., per-feature data sharing toggles) that allow Customer to operationalize this provision.

Note on Unlearning. In practice, it will be extremely difficult, and likely impossible, to remove or change data that has been integrated into an AI model. As such, Customer's counsel should counsel business partners that data initially placed into AI services should be considered there for good.

Section 3: Use and Ownership of Internal Weights and Biases of the AI Model

This Section addresses the legal and practical implications of ownership and access to the internal weights and biases of the AI Model. Such ownership matters for auditability, bias mitigation, transparency, regulatory compliance, and long-term Customer control. This Section operates alongside Section 4, which governs the license to the AI Model itself.

DRAFTER'S NOTE

Why This Section Should Be Included. Use and ownership of the internal weights and biases of the AI Model is essential for several reasons:

Mitigation of Hidden Risks. AI models often act as "black boxes," making it difficult to identify where bias occurs. Some level of access to or auditability of model weights enables Customer to assess and mitigate discriminatory outcomes, such as skewed hiring algorithms or biased customer segmentation.

Preventing Self-Reinforcing Feedback Loops. AI models can reinforce existing inequities embedded in historical data, perpetuating unfair patterns over time. Auditability supports continuous testing with diverse inputs to prevent training-serving skew, where the model fails in real-world scenarios.

Accountability and Trust. If Provider keeps the model fully opaque, Customer cannot fully explain how decisions are made. Greater transparency into model weighting supports stakeholder trust and helps Customer satisfy regulatory requirements, especially in critical sectors like healthcare, financial services, or HR.

Long-Term Accuracy and Control. Models need continual updating to reflect new data and evolving social norms. Ownership or licensed access to the model supports Customer in maintaining or refining performance without relying solely on Provider.

Cross-Reference. The ownership framework in this Section operates within the licensing structure of Section 4 and depends substantially on the delivery mode (SaaS, API, or on-premises).

Section 3.1: Provider Ownership of AI Model; Customer Ownership of Customer-Specific Customizations and Outputs

Customer shall own all Outputs, Deliverables, and Customer-Specific Customizations resulting from the AI Services provided to Customer under this Agreement, as further set forth in Sections 4 and 5. Provider shall retain ownership of the AI Model, the Provider's Proprietary Model, and any underlying Training Data, weights, parameters, and architectures developed using Provider's own resources without Customer Data, subject to the rights granted to Customer herein and the restrictions on use of Customer Data and Customer-Derived Improvements set forth in Sections 2 and 4.

DRAFTER'S NOTE

The Ideal Opening Position vs. Practical Reality. From Customer's perspective, the strongest position is ownership of the AI Model itself, including all weights and biases. In practice, this is rarely achievable in B2B AI procurement: Provider's entire business model is typically built on retaining ownership and reusing the Model across customers, outside of bespoke projects where Customer is specifically procuring, and willing to pay a premium for, a customized model. Counsel should pursue this maximalist position only where: (a) Provider is developing a model specifically for Customer; (b) Customer is funding substantially all development costs; and/or (c) the engagement is structured as a custom development project rather than a SaaS subscription.

Fallback Strategy. Where outright Model ownership is not achievable, the typical fallback ladder is: (i) exclusive license to the Model; (ii) non-exclusive license with field-of-use restrictions or competitor exclusion; (iii) non-exclusive license with audit and transparency rights; and (iv) standard non-exclusive license. The fallback positions below address several of these alternatives. Counsel should also negotiate strong portability and termination rights under Section 4.2.3 as a substitute for ownership.

FALLBACK 1 — EXCLUSIVE LICENSE TO CUSTOMER-SPECIFIC VARIANTS

Where Provider declines ownership transfer but the engagement includes meaningful Customer-Specific Customization:

"Provider grants Customer an exclusive, perpetual, irrevocable, royalty-free license to use, reproduce, and exploit any Customer-Specific Fine-Tuning of the AI Model performed using Customer Data. Provider shall not provide, license, or otherwise make available any such Customer-Specific Fine-Tuning to any other customer or third party."

Risk: Provider retains title to the underlying base Model. Customer's exclusivity attaches only to fine-tuned variants. This may be acceptable where the value to Customer lies in the customization, but does not address concerns about Provider improving its general Model based on insights from Customer's deployment (addressed in Section 4.3).

FALLBACK 2 — FIELD-OF-USE RESTRICTION

"Provider may use Customer-Specific Fine-Tuning of the AI Model only outside of Customer's field of use, defined as [insert industry, geographic, or use-case scope]. Provider shall not provide, license, or make available any Customer-Specific Fine-Tuning, or any AI Model trained or fine-tuned using Customer Data, to any direct competitor of Customer."

Risk: Enforcement of field-of-use restrictions can be challenging in practice, particularly where Provider serves many customers and the boundary between Customer's field and adjacent fields is contested. The definition of "direct competitor" may also require negotiation. Counsel should consider pairing this with representation and warranty obligations under Section 4.4.

Section 4: License to the AI Model

This Section addresses the license rights to the AI Model itself, distinct from rights in Customer Data (see Sections 1–2), Customer rights in weights and biases (Section 3), or rights in Outputs (see Section 5). It covers: (a) the rights granted to Customer based on the delivery mode (SaaS, API, or on-premises); (b) general restrictions on Customer's use; (c) disclosure and pass-through obligations regarding any Base Foundation Model; (d) ownership of the Provider's Proprietary Model; (e) ownership and treatment of Customer-Specific Customization; (f) restrictions on Provider's use of Customer Data in model development; and (g) Provider's representations, warranties, and covenants.

DRAFTER'S NOTE

Perspective. This Section is drafted from the Customer's perspective. Opening positions favor Customer protection. Fallback positions represent more balanced language a sophisticated Provider might accept. Risks introduced by fallback positions are identified explicitly.

Delivery Mode Matters. Section 4.1 is structured around three delivery modes: SaaS / Hosted AI Services, API Access, and Embedded / On-Premises Deployment. The legal analysis, particularly around whether a model license is even required, differs significantly between these modes. Counsel should confirm which mode applies before applying the rest of Section 4.

Section 4.1: Rights Based on Delivery Modes

Section 4.1.1: SaaS / Hosted AI Services

Where Provider delivers the AI Services as a hosted, cloud-based, or software-as-a-service offering in which Customer accesses the AI Model remotely and does not receive a copy of the model ("Hosted AI Services"), Customer does not receive a license to the AI Model itself. Instead, Provider grants Customer a limited, non-exclusive, non-transferable right to access and use the Hosted AI Services during the Term, solely for Customer's internal business purposes as set forth in the applicable Order Form or Statement of Work (the "Service Access Right").

In Hosted AI Services arrangements, the AI Model remains at all times within Provider's infrastructure and under Provider's control. Provider's obligations with respect to the AI Model in this context are service obligations, not license obligations, and include the commitments set forth in Sections 4.3 and 4.4 below.

DRAFTER'S NOTE

Why No Model License Is Required in SaaS. In a pure SaaS deployment, Customer sends Inputs to Provider's infrastructure and receives Outputs back. Customer does not reproduce, execute, or possess the AI Model — and therefore does not engage in any act that requires a copyright license to the model. The traditional software license grant is a holdover from on-premises deployment that does not map cleanly onto this architecture. Framing Customer's

rights as a “Service Access Right” rather than a model license more accurately reflects the legal and technical reality.

Model Extraction Risk. Although Customer does not possess the AI Model in a SaaS arrangement, Customer could in theory attempt to extract a functional equivalent through systematic querying (a practice known as model extraction or distillation). This risk is addressed in Section 4.1.4 (General Restrictions), which applies across all delivery modes.

Internal Business Purposes. “Internal business purposes” should be defined or clarified in a Statement of Work or Exhibit to prevent Provider from arguing that Customer’s use of Outputs in customer-facing products exceeds the Service Access Right.

FALLBACK — BROADER PERMITTED USE

If Provider resists the “internal business purposes” limitation:

“Provider grants Customer a limited, non-exclusive, non-transferable Service Access Right to access and use the Hosted AI Services for the purposes set forth in the applicable Order Form or Statement of Work.”

Risk: Shifts scope definition to the Order Form, which may receive less careful legal review. Ensure Order Forms are reviewed by counsel and use cases are described with specificity.

Section 4.1.2: API Access to the AI Model

Where Provider makes the AI Model available to Customer via application programming interface (API), software development kit (SDK), or similar programmatic access mechanism (“API Access”), Provider grants Customer a limited, non-exclusive, non-transferable, non-sublicensable license to use the API, SDK, and associated tooling and documentation solely to integrate the AI Model’s capabilities into Customer’s own applications and workflows, for Customer’s internal business purposes, during the Term (the “API License”).

The API License is a license to the API, SDK, and associated interface materials — not to the underlying AI Model itself, which remains hosted on Provider’s infrastructure. Customer’s right to make inference calls to the AI Model via the API is governed by the Service Access Right set forth in Section 4.1.1, which applies concurrently with the API License.

DRAFTER'S NOTE

API Access vs. Model Possession. Even in an API arrangement, Customer typically does not possess the model weights — API calls are routed to Provider’s hosted infrastructure. The API License therefore governs the client-side tooling (SDK, libraries, interface) while the Service Access Right from Section 4.1.1 governs the actual model inference. Counsel should confirm with Provider whether any model components are downloaded to Customer’s environment as part of API integration (e.g., embedding models, local inference caches), as these would require the on-premises license in Section 4.1.3.

Rate Limits and Usage Restrictions. API access agreements typically include usage-based rate limits and tiered pricing. Counsel should ensure that applicable rate limits, overage charges, and throttling policies are set forth in the Order Form with sufficient specificity to allow Customer to plan its usage.

Section 4.1.3: Embedded / On-Premises Model Deployment

Where Provider delivers the AI Model for installation and execution on Customer's own infrastructure or for embedding in Customer's own applications ("On-Premises Deployment"), Provider grants Customer a limited, non-exclusive, non-transferable, non-sublicensable license to: (a) install and execute the AI Model on Customer's designated systems solely for Customer's internal business purposes; and (b) incorporate the AI Model into Customer's internal applications and workflows, solely for Customer's internal business purposes, in each case during the Term and subject to the restrictions in Section 4.1.4 (the "On-Premises License").

The On-Premises License does not convey to Customer any ownership interest in the AI Model or any Model Intellectual Property. Customer's possession of the AI Model under an On-Premises License does not expand Customer's rights beyond those expressly granted herein.

DRAFTER'S NOTE

On-Premises License as True Software License. Unlike SaaS (Section 4.1.1) or API access (Section 4.1.2), an On-Premises Deployment involves Customer actually receiving and possessing a copy of the AI Model — typically as model weights and associated runtime software. This is the scenario where a traditional copyright license grant is legally necessary and appropriate. The restrictions in Section 4.1.4 are most legally significant in this context.

Escrow and Insolvency. Where Customer is operating the AI Model on its own infrastructure under an On-Premises License, Provider insolvency is a particularly acute risk — Customer may need updates, security patches, or support that a bankrupt Provider cannot provide. Counsel should negotiate model escrow arrangements and update obligations as part of any On-Premises Deployment.

EU AI Act Note. On-Premises Deployment of an AI Model may trigger obligations under the EU AI Act for Customer as a "deployer" of an AI system, including transparency, human oversight, and record-keeping obligations depending on the risk classification of the AI system. Counsel with EU nexus should assess these obligations before deployment.

Section 4.1.4: General Restrictions Applicable to All Delivery Modes

Regardless of the delivery mode, Customer shall not, and shall not permit any third party to:

- copy, reproduce, or distribute the AI Model or any portion thereof, except as strictly necessary for Customer's permitted use under Section 4.1.1, 4.1.2, or 4.1.3, as applicable;
- reverse engineer, decompile, disassemble, or otherwise attempt to derive the source code, model weights, Training Data, or underlying architecture of the AI Model, by any means, including without limitation through systematic or automated querying, adversarial probing, or model distillation techniques designed to replicate or approximate the AI Model's functionality;
- use the AI Model or AI Services to develop, train, fine-tune, or improve any competing AI model or service, except as expressly permitted under this Agreement;
- sublicense, sell, resell, transfer, assign, or otherwise dispose of the AI Model, the Service Access Right, the API License, or the On-Premises License, or any rights therein, without Provider's prior written consent;

- use the AI Model or AI Services in any manner that violates applicable law or the acceptable use policies disclosed by Provider to Customer in writing; or
- attempt to circumvent, disable, or interfere with any security, access control, or usage monitoring mechanisms implemented by Provider in connection with the AI Services.

DRAFTER'S NOTE

Model Distillation Across All Delivery Modes. Subsection (b) explicitly prohibits model extraction through systematic querying. This restriction is most practically significant in SaaS and API arrangements, where Customer lacks direct access to model weights but could in theory reconstruct a functional approximation through large-scale querying. The restriction closes the gap left by the absence of a traditional model license in SaaS deployments.

Competing AI Model Restriction. Subsection (c) prohibits using the AI Model to develop competing models. This is a Provider-favorable restriction that Providers often use as a starting point in negotiations. Counsel should confirm that this restriction does not inadvertently prevent Customer from: (a) using Outputs in Customer's own products (addressed in Sections 5–6); or (b) developing Customer's own internal AI tools that are not competitive with Provider's commercial offerings. Consider adding a carve-out for non-competitive internal tooling if Customer has active AI development programs. Alternatively, rather than relying on the relative and potentially evolving scope of "competing AI model or service" language, Customers can gain more certainty by defining what specific offerings amount to a competitive use.

FALLBACK — INTERNAL TOOLING CARVE-OUT

If Customer has active internal AI development programs:

"…except that Customer may develop internal AI tools and models for Customer's own internal business purposes, provided that such tools and models are not made commercially available to third parties and do not directly compete with Provider's commercially offered AI products or services."

Risk: Provider may resist this carve-out as difficult to police. The definition of "directly compete" may require further negotiation to avoid ambiguity.

Section 4.1.5: Base Foundation Model — Disclosure and Rights Sufficiency

To the extent the AI Model incorporates or is built upon a Base Foundation Model, the following obligations apply:

- **Universal Disclosure Obligation.** Provider shall disclose to Customer, prior to the Effective Date (or promptly upon Customer's written request at any time during the Term): (i) the identity of any Base Foundation Model(s) incorporated in the AI Model; (ii) a summary of the material use restrictions imposed by the licensor of such Base Foundation Model(s); and (iii) to the extent permitted by Provider's agreement with the Base Foundation Model licensor, a summary of the material terms of Provider's license to such Base Foundation Model(s). Provider shall promptly notify Customer of any material change to Provider's upstream licensing arrangements that could affect Customer's use of the AI Services.

- **Rights Sufficiency — Hosted AI Services and API Access.** For AI Services delivered as Hosted AI Services (Section 4.1.1) or via API Access (Section 4.1.2), Provider represents and warrants that it has obtained all rights, licenses, and permissions necessary to use the Base Foundation Model to provide the AI Services to Customer as contemplated by this Agreement, and that such use complies with the applicable license terms governing the Base Foundation Model.
- **Rights Sufficiency — On-Premises Deployment.** For AI Services delivered via On-Premises Deployment (Section 4.1.3), Provider represents and warrants that it has obtained all rights, licenses, and permissions necessary to: (i) use the Base Foundation Model in the AI Model as delivered to Customer; and (ii) grant Customer the On-Premises License set forth in Section 4.1.3 without violating the terms of any license governing the Base Foundation Model.
- **Customer Use Restrictions.** Customer's use of the AI Services shall be subject to any material use restrictions imposed by the licensor of the applicable Base Foundation Model, provided that Provider has disclosed such restrictions to Customer in writing prior to the Effective Date.

DRAFTER'S NOTE

Tiered Rights Sufficiency Rationale. The distinction between subsections (b) and (c) reflects an important legal point: in SaaS and API arrangements, no copyright license flows from Provider to Customer with respect to the Base Foundation Model — what matters is that Provider has the right to use the Base Foundation Model to deliver the service. In On-Premises Deployment, Provider must have sublicensable rights sufficient to convey the On-Premises License to Customer.

Open-Weight Models. If the Base Foundation Model is released under an open-weight license (e.g., Meta's Llama license, Mistral's Apache 2.0 license), confirm: (a) commercial use is permitted; (b) fine-tuning and redistribution of fine-tuned versions is permitted for the relevant delivery mode; and (c) attribution or notice requirements are satisfied. Open-weight licenses vary significantly in their restrictions.

Section 4.1.6: Provider's Proprietary Model — Ownership and License

As between the parties:

- Provider retains all right, title, and interest in and to the Provider's Proprietary Model, including all Model Intellectual Property therein, subject to the rights granted to Customer herein and the restrictions on use of Customer Data and Customer-Derived Improvements set forth in Sections 4.3 and 4.4.
- Provider's ownership of the Provider's Proprietary Model shall not include, and Provider hereby assigns to Customer, any improvements, enhancements, or modifications to the Provider's Proprietary Model that constitute Customer-Derived Improvements, as further described in Section 4.4.
- Nothing in this Agreement shall be construed to grant Customer any ownership rights in the Provider's Proprietary Model except as expressly set forth herein.

DRAFTER'S NOTE

Delivery-Mode Neutrality. This provision is intentionally delivery-mode-neutral with respect to ownership: Provider owns its Proprietary Model and Customer owns Customer-Derived Improvements regardless of whether the AI Services are delivered as SaaS, via API, or on-premises.

Practical Limitation in SaaS Context. While the assignment of Customer-Derived Improvements applies equally across all delivery modes, its practical value differs significantly. In a SaaS arrangement, Customer-Derived Improvements are embedded in Provider's hosted infrastructure and Customer cannot take possession of them. The assignment only has real practical value if paired with strong termination and portability rights under Section 4.2.3.

Assignment vs. License of Customer-Derived Improvements. The assignment in subsection (b) is a strong Customer-favorable position that Provider will likely resist, proposing instead an exclusive or non-exclusive license. See Section 4.4 and associated Fallback positions.

Section 4.2: Customer-Specific Customization

Section 4.2.1: Ownership of Customer-Specific Customizations

All Customer-Specific Customizations, and all intellectual property rights therein, shall be owned exclusively by Customer, subject to Section 4.2.2. Provider shall have no right to use, reproduce, distribute, disclose, or incorporate any Customer-Specific Customization into any product, service, or model, including without limitation the Provider's Proprietary Model or any model used by any other customer of Provider, except: (a) as strictly necessary to provide the AI Services to Customer under this Agreement; or (b) as expressly authorized in writing by Customer.

For the avoidance of doubt, Customer's ownership of Customer-Specific Customizations encompasses both Customer-Specific Configuration and Customer-Specific Fine-Tuning. The practical scope of Customer's ownership rights with respect to each sub-type, including limitations on possession of fine-tuned model weights in hosted deployments, is addressed in Section 4.2.3.

DRAFTER'S NOTE

Model Weights vs. Configuration. For Customer-Specific Fine-Tuning (adjustments to model weights), ownership by Customer may be technically complex in a hosted model environment. Customer's "ownership" may be a legal right without practical possession. For Customer-Specific Configuration (system instructions and parameter settings), ownership is more straightforward and possession is more achievable.

Work Made for Hire. If Provider performs the Customer-Specific Fine-Tuning as a service for Customer, consider whether the fine-tuning output could qualify as a "work made for hire" under 17 U.S.C. § 101. Given the unsettled nature of ML model training under work-for-hire doctrine, an express assignment clause is a more reliable mechanism.

FALLBACK — EXCLUSIVE LICENSE OR NON-EXCLUSIVE LICENSE WITH NON-USE COVENANT

If Provider resists full Customer ownership:

Fallback 1 — Exclusive License: "Provider grants Customer an exclusive, perpetual, irrevocable, royalty-free license to use, reproduce, modify, and exploit all Customer-Specific Customizations for any purpose, with the right to sublicense."

Risk (Fallback 1): Provider retains title, which creates risk in Provider bankruptcy. Exclusive licensee has stronger protections under 11 U.S.C. § 365(n) than a non-exclusive licensee, but ownership is stronger still.

Fallback 2 — Non-Exclusive License with Non-Use Covenant: "Provider grants Customer a non-exclusive, perpetual, irrevocable, royalty-free license to use Customer-Specific Customizations. Provider covenants it shall not use Customer-Specific Customizations for any other customer or incorporate them into any generally available product or service."

Risk (Fallback 2): Customer has only contractual protection, not an IP exclusivity right. Pair with audit rights (Section 11) and specific indemnification.

Section 4.2.2: Provider License Back

Customer hereby grants Provider a limited, non-exclusive, non-transferable license to access and use Customer-Specific Customizations solely to: (a) provide the AI Services to Customer under this Agreement; (b) maintain, support, and improve the AI Services as delivered to Customer; and (c) comply with applicable law. This license back shall not permit Provider to use Customer-Specific Customizations for any other purpose, including without limitation for the benefit of any other customer of Provider or for incorporation into the Provider's Proprietary Model.

DRAFTER'S NOTE

Why a License Back Is Necessary. If Customer owns Customer-Specific Customizations, Provider technically needs a license to use them in order to provide the AI Services. This provision is a functional necessity, not a concession.

Section 4.2.3: Portability and Termination Rights

Upon the expiration or termination of this Agreement, or upon Customer's written request at any time during the Term:

- Provider shall promptly deliver to Customer, in a reasonably accessible and documented format: (i) all Customer-Specific Configuration, including system prompts, parameter settings, and configuration files specific to Customer; and (ii) to the extent technically feasible, a complete snapshot or export of any Customer-Specific Fine-Tuning, including model weight differentials or adapter layers specific to Customer's fine-tuning;
- Provider shall provide reasonable transition assistance to Customer for a period of [ninety (90)] days following termination, at Customer's request, to facilitate Customer's migration to an alternative AI service provider;
- Provider shall certify in writing to Customer, within [thirty (30)] days following termination, that all copies of Customer-Specific Customizations (other than those delivered to Customer) have been destroyed or rendered permanently inaccessible, and shall provide a signed attestation of destruction upon Customer's request; and
- Provider's obligations under this Section 4.2.3 shall survive expiration or termination of this Agreement.

DRAFTER'S NOTE

Hosted Model Limitation. In a hosted/SaaS model, Provider will often resist delivering fine-tuned model weights. The phrase "to the extent technically feasible" is a practical concession. Customer's true protection is: (a) ensuring system configurations are fully portable; and (b) the certified destruction obligation.

Bankruptcy Intersection. If Provider enters bankruptcy before fulfilling termination obligations, Customer's ability to receive Customer-Specific Customizations will depend on bankruptcy treatment of the Agreement.

Transition Assistance Period. The [ninety (90)]-day transition period is a Customer-favorable starting position. Provider may seek to limit this to [thirty (30)] days. Consider

specifying that transition assistance is provided at no additional charge and define the scope of assistance.

FALLBACK — CONFIGURATION PORTABILITY ONLY

If Provider refuses to commit to delivery of fine-tuned model weights:

"Upon termination, Provider shall deliver to Customer all Customer-Specific Configuration in a documented, machine-readable format. Provider shall provide [sixty (60)] days of transition assistance at Customer's request. Provider represents and warrants that it shall not continue to use or benefit from any Customer-Specific Fine-Tuning following the date of certified destruction under Section 4.2.3(c)."

Risk: Customer loses practical access to fine-tuned weight differentials and will need to re-fine-tune with an alternative provider using its own data.

Section 4.3: Restrictions on Use of Customer Data in Model Development

Section 4.3.1: Prohibition on Use for General Model Improvement

Provider shall not use Customer Data, Customer-Specific Customizations, or any Customer-Derived Improvements, directly or indirectly, to:

- Train, fine-tune, retrain, update, or otherwise modify any AI Model that is used by, or made available to, any customer of Provider other than Customer;
- Improve the Provider's Proprietary Model or any Base Foundation Model in a manner that benefits any customer of Provider other than Customer; or
- Develop, enhance, or optimize any AI product or service offered by Provider to third parties.

DRAFTER'S NOTE

The "Indirect" Problem. The word "indirectly" is intentional. The hardest problem is not direct copying of Customer Data into training datasets — it is that Provider's engineers may develop insights from observing Customer's deployment that they apply to the general model. This "insight laundering" is difficult to detect and prove. Audit rights (Section 11) and indemnification are important supplementary protections.

Cross-Reference. Confirm alignment with Sections 2.1 and 2.2 on use of Customer Data. The prohibition in this Section 4.3.1 is focused on model development; Section 2 may address a broader range of data use restrictions.

Section 4.3.2: Permitted Use of Aggregated and Anonymized Data

Notwithstanding Section 4.3.1, Provider may use aggregated and anonymized metadata derived from Customer's use of the AI Services for the sole purpose of generally improving the reliability, security, and performance of Provider's AI Services, subject to the following conditions:

- Such metadata must be aggregated across multiple customers and must not be capable of identifying Customer or any Customer-specific patterns, use cases, configurations, or competitive information;
- Anonymization must be performed using industry-standard techniques that render re-identification technically infeasible, and Provider shall not attempt to re-identify Customer from such aggregated data;
- Such metadata shall not include any Customer Data, Customer-Specific Customizations, or any information derived primarily from Customer's Inputs or Outputs; and
- Provider's use of such metadata shall not confer, directly or indirectly, any competitive advantage on any other specific customer of Provider derived from Customer's use of the AI Services.

DRAFTER'S NOTE

Defining "Aggregated and Anonymized." Provider will want a broad definition; Customer will want a narrow one. Consider specifying a minimum aggregation threshold (e.g., data

must reflect usage by at least [X] customers) and referencing a recognized anonymization standard.

Section 4.4: Provider Representations, Warranties, and Covenants

Provider represents, warrants, and covenants to Customer that, as of the Effective Date and throughout the Term:

(1) **Authority and Rights.** Provider has the full right, power, and authority to grant the Service Access Right, the API License, the On-Premises License, and all other rights set forth in this Section 4, and such grant does not violate any agreement between Provider and any third party, including without limitation any agreement governing a Base Foundation Model;

(2) **No Conflicting Encumbrances.** The AI Model is free and clear of any liens, encumbrances, or restrictions that would prevent or limit Customer's exercise of the rights granted herein;

(3) **Compliance with Base Foundation Model Terms.** Provider's use of any Base Foundation Model in the provision of AI Services to Customer complies with the applicable license terms governing such Base Foundation Model, and Customer's use of the AI Model as permitted herein will not breach any such terms;

(4) **No Customer-Derived Improvements in General Model.** Provider has not incorporated, and shall not incorporate, any Customer-Derived Improvements into any AI Model used by or available to any customer of Provider other than Customer, or into the Provider's Proprietary Model as generally offered;

(5) **Customer-Specific Customizations Not Shared.** Provider has not used, and shall not use, any Customer-Specific Customization for the benefit of any other customer of Provider or in any generally available product or service; and

(6) **Notification of Material Changes.** Provider shall promptly notify Customer of any material change to the underlying AI Model, including any change in the Base Foundation Model, any significant retraining or architectural change to the Provider's Proprietary Model, or any change in Provider's upstream licensing arrangements that could materially affect Customer's rights under this Agreement.

DRAFTER'S NOTE

Representations as Enforcement Mechanism. Many restrictions imposed on Provider are difficult to verify technically. These representations and warranties serve as the primary enforcement mechanism. A breach gives Customer a breach-of-contract claim and may support a misappropriation of trade secrets claim under the Defend Trade Secrets Act (18 U.S.C. § 1836 et seq.) or applicable state law. Pair with: (a) specific indemnification obligations; (b) audit rights (Section 11); and (c) a right to terminate for material breach.

Warranty on Model Changes. Customer should have the right to receive advance notice of material changes and, if such changes adversely affect Customer: (a) reject the update and continue on the prior version for a reasonable period; or (b) terminate for convenience without penalty.

Section 5: Ownership and Exclusivity of Outputs from the AI

This Section addresses ownership of Outputs and Deliverables generated by the AI Services. The opening position vests ownership in Customer. This Section operates alongside Section 6 (Use Restrictions on Outputs) and Section 7 (Liability for Use of Outputs).

Section 5.1: Customer Ownership of Outputs and Deliverables

As between the parties, Customer shall own all right, title, and interest, including all intellectual property rights, in and to all Outputs and Deliverables generated by the AI Services in response to Customer's Inputs. Provider hereby assigns to Customer all such right, title, and interest, including all moral rights to the extent assignable under applicable law, agrees to execute such further documents as may be reasonably necessary to perfect such assignment, and waives, and agrees not to assert, any unassignable moral rights, to the extent permitted by applicable law.

Provider shall not use any Outputs or Deliverables for any purpose other than (a) providing the AI Services to Customer under this Agreement; or (b) as expressly authorized in writing by Customer. Provider shall not retain any rights in Outputs or Deliverables except as expressly set forth in this Agreement.

DRAFTER'S NOTE

Statutory Limitations on Output Ownership. Counsel should be aware that under current U.S. Copyright Office guidance, AI-generated content lacking sufficient human authorship may not be protectable by copyright at all. The assignment language above is therefore most effective with respect to: (a) Outputs that incorporate Customer's creative input through prompts, selections, and curation; (b) trade secret rights in the Outputs (independent of copyright); and (c) contractual rights as between Provider and Customer (regardless of underlying IP protectability). Counsel should not assume the assignment creates copyright ownership where copyright protection itself is unavailable.

Provider Use of Outputs. The prohibition on Provider's use of Outputs is important because Outputs may contain sensitive information derived from Customer Inputs. Provider may push back, particularly seeking to use Outputs for product improvement or training. Counsel should resist this on the same grounds as Section 2.1's restrictions on Customer Data: once Outputs enter a training pipeline, they may be irrecoverable.

Cross-Reference. Confirm consistency with any revisions to Section 6 (Use Restrictions) and Section 7 (Liability). Section 6 limits Customer's ability to use Outputs for certain prohibited purposes; that restriction is independent of and does not undermine Customer's ownership.

FALLBACK — JOINT OWNERSHIP OR LICENSE

Where Provider resists full Customer ownership:

Fallback 1 — Customer License: "As between the parties, Customer shall have an exclusive, perpetual, irrevocable, royalty-free, sublicensable, worldwide license to use, reproduce, modify, distribute, and otherwise exploit all Outputs and Deliverables for any purpose."

Risk (Fallback 1): Provider retains nominal title. Exclusivity provides most practical protections, but Customer cannot bring infringement claims as an owner and may face issues in Provider bankruptcy.

Fallback 2 — Non-Exclusive License with Non-Use Covenant: "Customer shall have a non-exclusive, perpetual, irrevocable, royalty-free, sublicensable license to all Outputs and Deliverables. Provider covenants that it shall not use Outputs or Deliverables generated for Customer for the benefit of any other customer or for the training or improvement of any AI Model."

Risk (Fallback 2): Customer has only contractual protection, not exclusivity. Breach detection may be difficult. Pair with strong audit rights under Section 11 and representation and warranty obligations.

Section 5.2: Provider Limited License to Operate the AI Services

Customer hereby grants Provider a limited, non-exclusive, non-transferable license to use Outputs and Deliverables solely to the extent necessary to: (a) provide the AI Services to Customer under this Agreement; (b) generate, process, and deliver the Outputs to Customer; and (c) comply with applicable law. This license does not permit Provider to use Outputs or Deliverables for any other purpose, including without limitation for the benefit of any other customer of Provider, for training or improving any AI Model, or for incorporation into any generally available product or service.

DRAFTER'S NOTE

Why a License Back Is Necessary. If Customer owns the Outputs, Provider technically requires a license to use them in connection with operating the AI Services (e.g., to deliver them back to Customer, log them for support purposes, etc.). This provision is a functional necessity, not a concession. Counsel should ensure the scope is tightly defined.

Section 6: Use Restrictions on Outputs

This Section addresses restrictions on Customer's use of Outputs. While Customer owns the Outputs under Section 5, certain categories of use are prohibited or restricted to protect Provider, third parties, and the public. The opening position below is calibrated to protect Customer from overreach while accepting reasonable restrictions on harmful or illegal use cases.

DRAFTER'S NOTE

Why This Section Should Be Included. Carefully defining the scope of permissible Output use is important because Outputs from AI models can be misused in ways that create liability for both Customer and Provider, and because the reliability and accuracy of AI Outputs in many contexts remains uncertain. At the same time, customers reasonably expect some baseline level of utility from the AI Services.

Two Versions Available. The opening position below adopts a concise approach with Customer-favorable refinements. A more Provider-favorable verbose alternative is provided as a fallback for transactions where Provider has substantial leverage or where the risk profile justifies broader restrictions.

Customizing the Prohibited Use List. The bulleted list of prohibited purposes is illustrative and should be edited based on the specific AI Services. For example: voice synthesis system providers should add explicit prohibitions against voice authentication fraud; image generation system providers should prohibit generation of deepfakes; and AI system providers used in regulated industries (healthcare, finance, employment) should clarify expected limitations based on applicable law.

Technical Backstops. Contract language alone cannot prevent misuse of AI Services. Where the risk profile warrants, Provider should also implement technical safeguards (input/output filtering, content moderation, rate limiting) to limit the likelihood of offensive, illegal, or dangerous content. Customer may wish to negotiate disclosure of these safeguards or contractual commitments to maintain them.

High-Risk Use Cases. Where the AI Services are intended for a particular high-risk purpose, it may be helpful to define that purpose and its expected limitations explicitly. For example, if the system is intended for use in cancer diagnosis support, the Agreement should specify that it is anticipated as a decision-support tool but that final diagnosis decisions remain the responsibility of a qualified healthcare provider.

Section 6.1: Acknowledgment of AI Limitations

Customer acknowledges that the AI Services operate probabilistically and may generate Outputs that are inaccurate, incomplete, biased, offensive, and/or otherwise unsuitable for the intended use. Customer is responsible for implementing safeguards, review, and/or validation appropriate to Customer's use case prior to relying on or externally using any Outputs. Outputs do not represent Provider's views or opinions.

Customer acknowledges that the AI Services may not be appropriate for use in connection with legal, medical, safety-critical, and/or regulatory decisions, and/or in any context where errors or omissions in Outputs could result in material harm. If Customer uses the AI Services in such contexts, Customer shall implement robust human review and independent verification of Outputs commensurate with the risk profile of the use case.

DRAFTER'S NOTE

Acknowledgments vs. Disclaimers. This Section is framed as Customer acknowledgments rather than Provider disclaimers. Acknowledgments hold the Customer responsible for making informed decisions about appropriate use, while disclaimers attempt to shift the legal risk to Customer. This section pairs with Customer's ownership of Outputs under Section 5: Customer both owns the Outputs and is responsible for evaluating their fitness for Customer's purposes.

High-Risk Use Cases — Contractual vs. Technical Allocation. If Customer's intended use is in a high-risk context (medical, legal, financial decisions, public safety, etc.), counsel should consider whether to: (a) require Provider to provide specific functionality (*e.g.*, confidence scores, uncertainty quantification) to support Customer's human review; (b) require Provider to meet specific performance standards (see Section 10 and the SLA Exhibit); and/or (c) prohibit use in certain contexts altogether. Pure acknowledgment language may not be sufficient in heavily regulated industries.

FALLBACK (VERBOSE / PROVIDER-FAVORABLE VERSION)

The following more verbose version is more protective of Provider and may be appropriate where Provider has substantial leverage or the AI Services involve unusually high-risk capabilities:

"The AI Services operate probabilistically and may generate Outputs that are inaccurate, incomplete, offensive, or unsuitable for the intended use. Customer is solely responsible for implementing appropriate safeguards, review, and validation prior to any reliance or external use of Outputs and agrees that it is solely responsible for evaluating all Outputs for accuracy and appropriateness. Customer further acknowledges that it should not rely upon Outputs for any purpose without careful manual review and cross-checking with other sources. Outputs should not be used in decision-making contexts where errors could reasonably be expected to result in material harm, including legal, medical, safety-critical, and regulatory compliance decisions, without comprehensive human review and independent verification."

Risk: The verbose version shifts more risk to Customer and may be unfavorable in dispute resolution. Customer should accept this language only after carefully evaluating whether the AI Services are likely to be relied upon in contexts where errors could cause material harm.

Section 6.2: Prohibited Uses

Customer shall not use the AI Services or Outputs for any of the following purposes:

- any illegal, fraudulent, deceptive, defamatory, and/or infringing purpose;
- operation of public safety, emergency response, and/or other services responsible for the immediate health or safety of individuals, except as expressly permitted in writing by Provider and supported by appropriate qualifications and safeguards;

- discriminating against any individual or group, and/or otherwise violating the civil or legal rights of any party, including in connection with employment, housing, credit, insurance, healthcare, or educational decisions;
- developing, training, fine-tuning, benchmarking, validating, or improving any competing or substitute artificial intelligence model or service, except as expressly permitted under this Agreement;
- creating pornographic, salacious, and/or gratuitously offensive material;
- generating content intended to impersonate identifiable individuals or entities, including deceptive synthetic media, except with the express, informed consent of the impersonated party (and/or, in the case of public figures, in clearly identified parody, commentary, and/or satire);
- extracting, inferring, reconstructing, and/or reverse engineering the AI Services or their underlying AI Model, architecture, Training Data, weights, or parameters, including through systematic querying or adversarial probing;
- attempting to circumvent any rate limits, access controls, security measures, and/or other restrictions implemented in connection with the AI Services; or
- any other use that Provider has identified as prohibited in Provider's acceptable use policy disclosed to Customer in writing prior to the Effective Date.

Customer shall not provide access to the AI Services to any third party except to its Authorized Users in accordance with this Agreement.

DRAFTER'S NOTE

Discrimination and Civil Rights. The prohibition against discrimination is particularly important if Outputs may be used in employment, housing, credit, insurance, healthcare, and/or educational decisions (that is, areas where federal and state civil rights laws impose specific protections). Counsel should be aware that simply prohibiting discrimination might not be sufficient if the AI Services themselves produce discriminatory outputs; this prohibition pairs with the accuracy and bias considerations in Section 10 and the auditability considerations in Section 3.

Synthetic Media and Deepfakes. The exception for impersonation with consent or for parody/commentary carves out legitimate, First Amendment-protected uses, such as creative, journalistic, and commentary uses.

Acceptable Use Policy Hook. The final bullet incorporates Provider's acceptable use policy by reference. Customer's counsel should require that any such policy be disclosed in writing prior to signing, and that material changes require Customer's consent or trigger a termination right. Without such guardrails, Provider could effectively amend the contract unilaterally by changing the acceptable use policy.

Section 7: Liability for Use of the Outputs

This Section allocates liability for Customer's use of Outputs. It addresses three distinct categories: (a) direct liability between the parties; (b) third-party liability for Customer's use; and (c) statutory and regulatory liability for unwanted outcomes. The opening position preserves meaningful Provider accountability while recognizing that Customer typically controls the use of Outputs and is best positioned to assess fitness for purpose.

Section 7.1: Direct Liability Between the Parties

Provider may modify the AI Services from time to time. Customer acknowledges that modifications (including ongoing model training, retraining, fine-tuning, and software updates) may cause the AI Services' performance to change over time. Provider shall provide reasonable advance notice of material changes pursuant to the notification of material changes covenant set forth in Section 4.4.

Subject to the warranties, indemnities, and service levels expressly set forth in this Agreement, Customer is responsible for its Inputs to the AI Services and its use of Outputs from the AI Services. Provider shall not be liable to Customer for the consequences of Customer's use of Outputs, including consequences arising from Customer's implementation of Outputs in Customer's products, services, decisions, or workflows, except to the extent such consequences arise from Provider's breach of this Agreement (including any breach of representations, warranties, or service level commitments).

DRAFTER'S NOTE

Why This Section Should Be Included. Providers might seek to place liability for the consequences of AI use entirely on the Customer on the theory that Providers cannot actively monitor each Customer's use. While a reasonable starting principle, a pure shift-of-all-risk approach may be unfavorable to Customer. The opening position preserves Customer responsibility for use decisions while keeping Provider on the hook for breaches of express commitments including warranties (Section 4.4), service levels (Section 10), and security obligations (Section 11).

Interaction with SLAs. Provider may not take responsibility for AI performance on the theory that AI Services operate probabilistically and their accuracy cannot be guaranteed. While that might be true, there is room for negotiation to allow Customer to obtain some level of Provider accountability. If Customer expects specific levels of accuracy, latency, throughput, or availability, those should be addressed in the SLA Exhibit (referenced by Section 10), with Provider remaining accountable for material breaches of those SLAs.

FALLBACK (PROVIDER-FAVORABLE DIRECT LIABILITY)

If Provider insists on broader disclaimer language:

"The AI Services may be modified by Provider at any time. No guarantee is provided that operation of the AI Services will stay the same. Customer acknowledges that modifications may cause AI Services performance to change. Customer maintains sole responsibility for its Inputs and use of Outputs. Provider shall not be responsible for Customer's reliance on Outputs."

Risk: This formulation arguably eliminates Provider accountability even for breaches of express warranties or service levels. Customer should consider pairing with strong express commitments elsewhere in the Agreement.

Section 7.2: Third-Party Liability for Customer's Use

Subject to Section 8 (Third-Party IP Infringement Relating to Training Data) and the warranties expressly set forth in this Agreement, Customer is responsible for ensuring that its Inputs and its use of Outputs do not infringe, misappropriate, and/or violate the rights of any third party. Customer represents and warrants that it has all rights, licenses, consents, and permissions necessary to provide its Inputs to the AI Services and to use the Outputs as Customer intends.

Customer shall indemnify and hold harmless Provider from and against any and all Losses arising out of or related to: (a) Customer's breach of this Agreement, including the prohibited uses set forth in Section 6.2; (b) any third-party claim arising out of Customer's Inputs (other than claims for which Provider is responsible under Section 8 or Section 4.4); and/or (c) any downstream use of Outputs by Customer's customers, clients, or other parties to whom Customer provides access to the Outputs, except to the extent such Losses arise from Provider's breach of this Agreement.

Except as expressly set forth in this Agreement, Provider is not obligated to perform any pre- or post-processing of Inputs or Outputs, including any processing that would render Provider responsible for Customer's use of the AI Services.

DRAFTER'S NOTE

Why This Section Should Be Included. AI Outputs can infringe third-party rights in ways that are difficult to detect or predict, particularly with generative AI. Both parties have legitimate interests in clearly allocating responsibility: Customer chooses what Inputs to submit and how to use Outputs; Provider chose what Training Data to use and how to design the AI Model. This Section addresses the former; Section 8 addresses the latter.

Reputational Harm. Customer's use of AI tools, even where legally permissible, may create reputational risk that cannot be fully allocated by contract. Counsel should consider whether the Agreement requires Provider to provide attribution, disclosure, or branding controls that allow Customer to manage public-facing AI use. This is particularly important where the AI Services are used in public-facing contexts (chatbots, customer service) or in sensitive subject matter areas.

Carve-Outs for Gross Negligence and Fraud. A common carve-out from broad customer indemnification obligations is gross negligence or fraud by Provider. The opening position above preserves this by limiting Customer indemnification to claims, "except to the extent

such Losses arise from Provider's breach." Counsel may also wish to add explicit gross negligence and fraud carve-outs depending on the broader Agreement's structure.

FALLBACK (PROVIDER-FAVORABLE THIRD-PARTY INDEMNIFICATION)

If Provider insists on broader Customer indemnification:

"By using the AI Services, Customer acknowledges that, absent fraud or gross negligence on the part of Provider, Customer is solely responsible for ensuring that Inputs and use of Outputs do not infringe others' rights. Customer agrees to indemnify and hold harmless Provider for any third-party claims arising out of Customer's use of the AI Services, including all allegations of infringement of copyrights, trademarks, trade dress, patents, trade secret misappropriation, fraud involving use of the AI Services, or the like. Provider provides no non-infringement warranties for Outputs. Customer is solely responsible for its Inputs and use of Outputs. No pre- or post-processing of Inputs or Outputs by Provider makes Provider responsible for Customer's use of the AI Services. Customer is also solely responsible for any downstream use of Outputs."

Risk: This broader formulation eliminates Provider's warranty obligations regarding Training Data and shifts all third-party infringement risk to Customer, even for infringement arising from Provider's Training Data choices. Customer should resist this unless Section 8 provides robust counterbalancing protections.

Section 7.3: Statutory and Regulatory Liability for Unwanted Outcomes

Customer is responsible for ensuring that its use of the AI Services complies with all applicable laws and regulations in the jurisdictions where Customer operates and where Outputs are used or relied upon. Customer accepts the risk that regulatory changes, modifications to the AI Services, and/or other events may cause Customer's use of the AI Services to become subject to new or different legal requirements. Customer shall bear the responsibility for and costs of compliance with applicable laws and regulations as they apply to Customer's use of the AI Services.

Where the AI Services are intended for use in a particular regulated context (such as financial services, healthcare, employment, insurance, and/or other highly regulated areas), Customer shall comply with the regulatory framework applicable to that context. Customer shall not use the AI Services in any jurisdiction where Provider has notified Customer in writing that such use is prohibited or restricted.

DRAFTER'S NOTE

Why This Section Should Be Included. When AI Services are used in highly regulated or controversial areas (finance, medicine, insurance, employment, etc.), the regulatory framework may impose specific obligations on the user of AI tools that Provider cannot satisfy on Customer's behalf. This Section makes clear that compliance with those obligations is Customer's responsibility. Where Customer's intended use is in such a regulated context, counsel should also consider negotiating specific Provider obligations (e.g., compliance certifications, documentation, audit support) under Section 11 or in a separate compliance addendum.

Jurisdictional Considerations. If applicable, the Agreement may identify specific jurisdictions where the AI Services are prohibited or restricted. For example, certain AI applications are restricted under the EU AI Act, China's Interim Measures for the Management of Generative Artificial Intelligence Services, and/or specific state laws (*e.g.*, New York City's algorithmic hiring restrictions). Counsel should ensure that such restrictions are disclosed in writing and that the Agreement does not require Customer to use the AI Services in jurisdictions where doing so would be unlawful.

Provider Modifications. The acknowledgment that modifications to the AI Services may change their regulatory profile is intentional. AI systems' behavior can shift materially through retraining, fine-tuning, and/or model updates, potentially affecting compliance posture. This is one reason Section 4.4 requires advance notice of material changes: it gives Customer the opportunity to reassess compliance before the change takes effect.

Section 8: Third-Party IP Infringement Relating to Training Data

This Section addresses Provider's obligations regarding third-party intellectual property rights potentially affected by the AI Services, with particular focus on Training Data used to develop the AI Model. It provides warranties of non-infringement, indemnification, and remedies in the event of infringement claims.

DRAFTER'S NOTE

Why This Section Should Be Included. Liability for use of copyrighted works and other intellectual property in AI Training Data is unsettled. Customers typically have no visibility or control over model training, including what data Provider used and how it was obtained. Standard non-infringement clauses in software agreements may not fully capture risks specific to AI transactions, including the risk that the AI Model itself was trained on infringing data.

Risks of Not Including This Section. Without robust protection in this area, Customer may face: (a) liability for infringing Outputs generated by the AI Model, even where Customer had no intent to infringe and took reasonable steps to avoid infringement; (b) third-party claims based on Provider's use of copyrighted materials in Training Data; (c) substantial costs to defend such claims even where Customer is ultimately not found liable; and (d) loss of access to the AI Services if Provider is found liable for infringement.

Pirated Corpora Risk. Following cases such as *Kadrey v. Meta Platforms, Inc*., No. 3:23-cv-03417-VC, 2025 WL 4123456 (N.D. Cal. June 25, 2025) and similar litigation, Customer should be aware that some AI models have been trained on pirated corpora (including shadow libraries such as Library Genesis and Anna's Archive). Customer should consider Provider warranties regarding the lawfulness of Training Data acquisition rather than relying solely on Customer indemnification of downstream use.

Provider Resistance. Providers may resist broad warranties regarding Training Data because: (a) they may not have fastidiously documented the provenance of all Training Data, particularly for older models; (b) Training Data for Base Foundation Models is often outside Provider's control; and (c) the law in this area remains in flux. Counsel should expect substantial negotiation around the scope of warranties and indemnification.

Section 8.1: Warranty of Non-Infringement

Provider represents and warrants to Customer that, as of the Effective Date and throughout the Term:

(7) **Training Data Rights.** Provider's acquisition and use of Training Data, and the AI Services and AI Model, do not and will not infringe, misappropriate, or violate any Third-Party Intellectual Property Rights.

(8) **Lawful Acquisition.** Provider has obtained, and will maintain throughout the Term, all licenses, consents, and permissions required to collect, process, and use all Training Data in connection with providing the AI Services and AI Model, including for design,

development, creation, training, fine-tuning, use, implementation, and provision of the AI Services. Provider has not used Training Data acquired from sources known by Provider, or that reasonably should have been known by Provider, to be unlawful or unauthorized.

(9) **Output Non-Infringement.** Subject to Customer's use of Outputs in accordance with this Agreement, the Outputs generated by the AI Services in response to Customer's Inputs do not and will not infringe, misappropriate, or violate any Third-Party Intellectual Property Rights, including through substantial similarity to or memorization of Third-Party Intellectual Property in the Training Data.

DRAFTER'S NOTE

Scope of Warranties. The warranties in this Section are intentionally broad. Provider may seek to narrow them by: (a) limiting the warranty to Provider's own Training Data (excluding Base Foundation Model Training Data); (b) qualifying the warranty by Provider's knowledge; (c) carving out Outputs that result from Customer Inputs designed to elicit infringing content; or (d) limiting Output non-infringement warranties to specific categories of IP (e.g., copyright but not trademark).

Memorization Risk. The reference to "memorization" in subsection (c) addresses the AI-specific risk that a model may reproduce verbatim or near-verbatim portions of Training Data in its Outputs. This risk is heightened where Training Data is repeated in the corpus or where the model is over-fit. Counsel should consider whether to require Provider to implement specific technical controls to mitigate memorization risk.

Section 8.2: Indemnification

Provider shall defend, indemnify, and hold harmless Customer, its affiliates, and their respective directors, officers, employees, and agents from and against any and all Losses arising out of or related to any claim, allegation, or finding that the AI Services, AI Model, Training Data, Outputs, or Deliverables (or Customer's use of any of the foregoing in accordance with this Agreement) infringe, misappropriate, or violate any Third-Party Intellectual Property Rights, including allegations based on substantial similarity to or memorization of Third-Party Intellectual Property in the Training Data.

Provider's indemnification obligations under this Section 8.2 shall not be limited by any general limitation of liability provisions in the Agreement, except as expressly set forth herein. Customer shall provide Provider with prompt written notice of any claim subject to indemnification, sole control of the defense and settlement (provided that no settlement shall impose any non-monetary obligation on Customer without Customer's prior written consent), and reasonable cooperation in the defense.

DRAFTER'S NOTE

Caps and Carve-Outs. Provider will typically seek to: (a) cap indemnification at a multiple of fees paid or some other ceiling; (b) carve out indemnification for Customer's use in combination with non-Provider products or services; (c) carve out indemnification for Customer's use of unsupported versions of the AI Services; and (d) carve out indemnification for Inputs provided by Customer. Counsel should resist caps on IP indemnification in

particular because such damages are inherently unpredictable and can be very large. Where caps are accepted, they should be substantially higher than general damages caps.

Control of Defense. Provider will typically demand sole control of the defense and settlement of indemnified claims. The opening position grants this, with a protection that Provider cannot impose non-monetary obligations on Customer without consent. Counsel should also consider whether to require Provider to consult with Customer on material defense decisions and whether to require Customer's consent for any settlement that admits liability or wrongdoing.

Section 8.3: Infringement Remedies

If the AI Services, AI Model, Training Data, or Outputs (or any component thereof) are, or in Provider's reasonable opinion are likely to be, the subject of an infringement claim or finding, Provider shall, at its own expense and in addition to any indemnification or other remedies provided by this Agreement, use commercially reasonable best efforts to:

- procure for Customer the right to continue using the AI Services, AI Model, and Outputs as contemplated by this Agreement;
- modify or replace the AI Services, AI Model, or Outputs to make them non-infringing, while providing at least substantially equivalent functionality and performance; or
- if neither of the foregoing is commercially reasonable, promptly refund to Customer all prepaid, unused fees and any fees paid for any period during which Customer cannot use the AI Services. Customer's acceptance of such refund shall not be deemed a waiver of any other rights or remedies under this Agreement or applicable law.

DRAFTER'S NOTE

Order of Remedies. The order of remedies above (procure rights, modify, refund) is standard but the language matters. "Commercially reasonable best efforts" is a stronger formulation than "commercially reasonable efforts" and signals that Provider must genuinely attempt the first two options before falling back to refund. Provider may push to weaken this to a simple "commercially reasonable efforts" standard.

Refund Insufficiency. A refund of fees paid is often inadequate compensation for the disruption caused by losing access to the AI Services. Customer may have built systems around the AI Services, integrated them with Customer Data, and incurred substantial sunk costs. Counsel should consider whether to require additional remedies (transition assistance, cooperation with alternative providers, etc.) in the event of a refund-only outcome.

Non-Waiver Clause. The non-waiver language at the end of subsection (c) is important: without it, Provider might argue that Customer's acceptance of a refund constitutes acceptance of Provider's remedies and waiver of indemnification or other claims. The clause preserves Customer's rights.

Section 9: Improvements, Suggestions, and Feedback to the Provider

This Section addresses ownership and use of Feedback provided by Customer to Provider relating to the AI Services. The opening position protects Customer's underlying intellectual property in Customer-originated materials while permitting Provider to use Feedback to improve the AI Services.

Section 9.1: Customer Ownership of Confidential Information and Materials

As between the parties, Customer reserves and retains all right, title, and interest, including all intellectual property rights, in and to: (a) Customer's Confidential Information; (b) any proprietary know-how of Customer; (c) Customer Data and Inputs; and (d) any other text, photos, graphics, recordings, documentation, images, drawings, designs, plans, artwork, models, specifications, processes, software, algorithms, or other work product (including all related drafts and preparatory materials) created by Customer using the AI Services, excluding any Outputs or Deliverables that are otherwise allocated under this Agreement (collectively, "Customer Materials," as defined in the Definitions section). At no time will Provider acquire or retain any title to or ownership of any Customer Materials.

DRAFTER'S NOTE

Why This Provision Matters. Customer Materials is a defined category that captures Customer's underlying IP — distinct from Outputs (Section 5), Customer-Specific Customizations (Section 4), and Customer Data (Sections 1 and 2). The intent is to ensure that Customer's ownership of its pre-existing and independently created materials is not eroded by virtue of using the AI Services. This is particularly important where Customer uses the AI Services to refine, format, or process pre-existing Customer IP.

Section 9.2: Feedback

Customer or its Authorized Users may, from time to time, provide Feedback to Provider relating to the AI Services, AI Model, or Model Documentation. To the extent Customer provides such Feedback, and subject to Customer's ownership of Customer Materials and Customer Confidential Information under Section 9.1, Customer grants Provider a non-exclusive, perpetual, irrevocable, royalty-free, worldwide license to use such Feedback to improve, enhance, modify, and develop the AI Services and Provider's other products and services.

Notwithstanding the foregoing, this license does not grant Provider any rights in: (a) Customer Data, Inputs, Outputs, Deliverables, or Customer-Specific Customizations (which are addressed in Sections 1–5 and remain subject to the restrictions therein); (b) any Customer-Derived Improvements (which are addressed in Section 4); or (c) any Customer Confidential Information embedded in or referenced by the Feedback. Provider shall not use Feedback in any manner that discloses Customer's identity, business operations, or use cases without Customer's prior written consent.

DRAFTER'S NOTE

Feedback vs. Customer-Derived Improvements. These two categories are intentionally distinct. Feedback is what Customer says to Provider, and focused on suggestions, ideas, recommendations for Provider's AI Services. Customer-Derived Improvements are technical improvements to the AI Model derived from Customer Data, configuration, or usage — which are addressed under Section 4 and assigned to Customer (or licensed exclusively to Customer in the relevant fallback). The carve-out language in Section 9.2 ensures that the Feedback license does not inadvertently capture Customer-Derived Improvements.

Confidentiality Tension. Feedback inevitably contains information about Customer's use of the AI Services, which may itself be sensitive. The final sentence of Section 9.2 protects Customer's identity, business operations, and use cases from disclosure. Counsel should also confirm that the broader Agreement's confidentiality provisions apply to Feedback.

Cooperation Requirement. Provider may seek to add a cooperation requirement obligating Customer to execute documents necessary to perfect the Feedback license. This is generally unobjectionable so long as it does not require Customer to assign or transfer any Customer Materials or Customer-Derived Improvements.

FALLBACK (PROVIDER-FAVORABLE FEEDBACK TREATMENT)

Where Provider has substantial leverage and insists on broader Feedback rights:

"Subject to the rights and licenses granted herein by Provider, Provider, its affiliates, and its licensors reserve and retain all rights, title, and interest (including all intellectual property rights) in and to: (a) the AI Services, AI Model, and Documentation; (b) Provider's Confidential Information; (c) all modifications, enhancements, or improvements to the foregoing developed by or for Provider using Provider's own resources and not derived from Customer Data, Customer-Specific Customizations, or Customer's use of the AI Services (the "Provider Improvements"); and (d) any Feedback. Provider shall be free to use, disclose, reproduce, license, and otherwise distribute the Provider Improvements and Feedback without obligation or restriction; provided that, for the avoidance of doubt, Customer-Derived Improvements (as defined herein) are expressly excluded from the scope of "Provider Improvements" and remain subject to Section 4."

Risk: This formulation grants Provider ownership of all Improvements without regard to whether they qualify as Customer-Derived Improvements under Section 4, creating tension with Section 4.1.6 and 4.4. If Customer accepts this fallback, the definition of "Improvements" must be carefully drafted to exclude Customer-Derived Improvements, or the fallback should expressly state that Section 4 controls in the event of conflict.

Alternative positions contemplate exclusive rights regarding use of Feedback, such as in certain fields or applications which allows Provider the benefit of the Feedback while limiting Competitor's competition.

Section 10: Possible Future Issues

This Section addresses accuracy requirements, AI-specific service level commitments, and the right to pin to specific Model versions. The detailed service level metrics are set forth in the SLA Exhibit referenced below; this Section provides the operative warranty and remedy framework.

DRAFTER'S NOTE

Why This Section Should Be Included. Reliance on AI is inherently risky for Customer given hallucinations, biases, and other inaccuracies that may not be evident at the time of use. At the same time, vendors typically push for broad disclaimers of accuracy and shift performance responsibility to Customer. This Section establishes a baseline of Provider accountability for AI Services performance, with specifics set forth in the SLA Exhibit.

Calibrating to the Use Case. The appropriate level of accuracy and reliability commitment depends heavily on the use case. For low-stakes applications (drafting assistance, content brainstorming), aspirational language may suffice. For high-stakes applications (medical decision support, financial analysis, hiring decisions), specific quantitative commitments and remedies for breach are essential. Counsel should ensure that the SLA Exhibit reflects the actual risk profile of Customer's use.

Model Version Pinning. AI models can change materially through retraining or fine-tuning, sometimes in ways that degrade performance for specific use cases. The right to pin to a specific Model version protects Customer from unwanted regressions, but Provider may resist on the grounds that maintaining older versions is costly. Where pinning is critical, counsel should confirm Provider's technical capability to support specific versions and the cost (if any) for doing so.

Section 10.1: Performance and Accuracy Warranty

Provider represents and warrants that:

(10) **Material Accuracy.** The AI Services shall perform accurately in all material respects, and the Outputs shall be materially accurate, in accordance with the Documentation and the Accuracy Targets set forth in the SLA Exhibit.

(11) **Service Level Compliance.** The AI Services shall meet or exceed the Accuracy Targets, Latency, Throughput, Safety, and Availability service levels set forth in the SLA Exhibit.

(12) **Hallucination and Refusal Rates.** For designated tasks specified in the SLA Exhibit, the Hallucination Rate shall not exceed the maximum thresholds set forth therein, and the Refusal Rate for permitted tasks shall not exceed the maximum thresholds set forth therein.

(13) **No Material Degradation.** Provider shall not implement changes to the AI Services that cause a material degradation in any service level or Accuracy Target. Where Provider proposes any change reasonably likely to cause such degradation, Provider shall provide advance written notice to Customer and offer Customer the

opportunity to: (i) reject the change and continue on the prior version; or (ii) terminate the affected portion of the Agreement without penalty.

DRAFTER'S NOTE

Defining Accuracy. "Accuracy" for AI Services is not a single concept. Counsel should define accuracy in a manner appropriate to the use case, which may include: (a) factual accuracy (correctness of information conveyed); (b) task accuracy (correctness of completed tasks); (c) format accuracy (Output conforms to specified format); (d) classification accuracy (correctness of categorization). The SLA Exhibit should specify which forms of accuracy are measured and how.

Vendor Resistance. Providers will resist quantitative accuracy commitments on the grounds that: (a) AI accuracy is inherently probabilistic; (b) measurement methodology may be disputed; and (c) accuracy depends on Customer Inputs over which Provider has no control. Customer should respond by: (a) accepting that accuracy commitments are statistical, not absolute; (b) specifying clear measurement methodology in the SLA Exhibit; and (c) limiting accuracy commitments to designated task categories where measurement is feasible.

Section 10.2: Model Version Pinning

Customer shall have the right to pin its use of the AI Services to specific AI Model versions or weight sets for at least [twelve (12)] months from the general availability of such version (the "Pinning Period"). During the Pinning Period, Provider shall maintain the pinned Model version and shall not modify, retrain, or replace it in a manner that materially affects Customer's use. After the Pinning Period, Provider may discontinue support for the pinned Model version with [ninety (90)] days advance written notice.

DRAFTER'S NOTE

Why Pinning Matters. Even small changes to an AI Model can produce material changes in Outputs. For Customers that have validated the AI Services for specific use cases, an unannounced model update can break dependent systems, invalidate compliance certifications, or change Output quality. Pinning gives Customer the time to test and re-validate before adopting model changes.

Negotiating the Pinning Period. The 12-month default is a starting point. Customer may need longer (e.g., 24 months) for use cases requiring extensive validation, or shorter (e.g., 6 months) where the benefits of latest model improvements outweigh stability needs. Provider may push back on long pinning periods because maintaining old versions is costly. Where pinning is critical, consider whether to: (a) require Provider to maintain a published list of supported versions and EOL dates; (b) require Provider to provide migration tools to ease transition between versions; and (c) require Provider to commit to backward compatibility between versions.

Section 10.3: SLA Exhibit

The specific Accuracy Targets, Latency, Throughput, Safety, Availability, Hallucination Rate, Refusal Rate, and other quantitative service level commitments, together with measurement

methodology and remedies for breach, are set forth in [Exhibit A — AI Service Level Agreement] attached hereto and incorporated by reference.

DRAFTER'S NOTE

The SLA Exhibit. See Exhibit A at the end of this Toolkit for a template AI Service Level Agreement, including drafter guidance on each metric, typical ranges, and considerations for tailoring to specific use cases. Revise the listed service level commitments in Sections 10.1 and 10.3 to match those in the final version of Exhibit A.

Section 11: Security Issues to Consider

This Section addresses Provider's security obligations with respect to Customer Data and the AI Services. AI systems materially expand traditional security risk in ways that standard SaaS clauses do not fully capture, and this Section is structured to address both traditional security obligations and AI-specific risks.

DRAFTER'S NOTE

Why This Section Should Be Included. AI systems expand the traditional security risk profile in several important ways: (a) expanded attack surface, including new vectors such as APIs, prompts, model layers, embeddings, and fine-tuning pipelines beyond typical SaaS infrastructure; (b) non-obvious data leakage risks, including leakage through model outputs, embeddings, logs, and retraining artifacts even where raw data is nominally secured; (c) irreversibility of exposure, because Customer Data incorporated into a model (whether intentionally or not) may be impossible to fully remove; (d) cross-customer contamination risk, where one customer's data can influence another's outputs without strict controls; and (e) regulatory and contractual dependency, because Customer remains accountable for downstream security failures even when caused by Provider.

Risks of Not Including This Section. Without robust security provisions, Customer is exposed to: data leakage through AI systems (e.g., Customer Data appearing in other customers' outputs); silent model training on Customer Data; inadequate breach response and delayed notification; security failures outside traditional scope (prompt injection, model inversion); regulatory non-compliance cascades (GDPR, CCPA); and inability to investigate incidents due to lack of logs or audit rights.

Vendor Negotiating Positions. Counsel should expect Provider to push back in the following ways: (a) substitute "commercially reasonable efforts" for specific obligations; (b) replace specific controls (MFA, encryption standards) with framework-based compliance only (e.g., "SOC 2 compliant"); (c) resist AI-specific commitments around training pipelines, model leakage, prompt injection, or model inversion; (d) replace audit rights with summary reports only and limit penetration testing visibility; (e) extend breach notification windows ("without undue delay" instead of 24–72 hours); (f) demand broad subprocessor discretion with minimal approval rights; and (g) decline guarantees on segregation in multi-tenant or shared model environments.

Section 11.1: Comprehensive Security Program

Provider represents and warrants that it has implemented and shall maintain throughout the Term and any period during which Provider possesses or accesses Customer Data, a comprehensive written information security program. This program shall include administrative, technical, and physical safeguards that are reasonably designed to:

- ensure the security, integrity, and confidentiality of Customer Data;
- protect against any anticipated or reasonably foreseeable threats or hazards to the security or integrity of Customer Data;

- protect against unauthorized access to, use, disclosure, alteration, or destruction of Customer Data;
- protect against AI-specific risks, including prompt injection, model inversion, model extraction, embedding leakage, and cross-customer data contamination; and
- comply with industry standards related to data security and privacy.

DRAFTER'S NOTE

AI-Specific Risks Explicitly Listed. Subsection (d) is intentionally specific. A general "security program" obligation may not encompass AI-specific attack vectors that are not part of traditional information security frameworks. The explicit enumeration ensures these risks fall within Provider's security obligations and provides a basis for breach claims if Provider fails to address them.

Section 11.2: Administrative Safeguards

Provider's administrative safeguards shall include, at a minimum:

- **Security Awareness and Training:** regular mandatory security awareness and data handling training for all personnel (including employees and contractors) authorized to access Customer Data, with AI-specific security training for personnel involved in operating, maintaining, or modifying the AI Services;
- **Designated Security Responsibility:** designation of a qualified individual(s) responsible for overseeing and managing Provider's information security program;
- **Risk Assessment:** periodic risk assessments of potential threats and vulnerabilities to Customer Data and the AI Services, including AI-specific risks, and implementation of measures to mitigate identified risks;
- **Security Event Response Plan:** a written Security Event response plan, regularly tested, designed to promptly identify, contain, investigate, eradicate, and recover from any compromised Customer Data or AI Security Incident;
- **Personnel Security:** to the extent permitted by law, procedures for screening personnel with access to Customer Data, commensurate with the sensitivity of the Customer Data, and ensuring personnel are bound by appropriate confidentiality obligations; and
- **Subprocessor Management:** procedures for assessing and managing the security practices of any Subprocessors permitted under this Agreement that may access Customer Data.

Section 11.3: Technical Safeguards

Provider's technical safeguards shall include, at a minimum:

- **Access Controls:** technical controls ensuring that access to Customer Data is strictly limited based on the principle of least privilege, including: (i) unique user identifiers for each individual; (ii) robust authentication mechanisms (including multi-factor authentication for administrative access); (iii) logical access controls based on job function and need-to-know; (iv) regular review and prompt revocation of access rights; and (v) logging and monitoring of access to systems containing Customer Data;

- **Encryption:** encryption of Customer Data at rest using industry-standard strong encryption algorithms (e.g., AES-256 or equivalent) and in transit across public or untrusted networks (e.g., TLS 1.2 or higher);
- **Network Security:** firewalls, intrusion detection and prevention systems (IDPS), and other network security measures to protect systems containing Customer Data;
- **Vulnerability Management:** processes for regularly identifying and remediating security vulnerabilities, including timely application of security patches and updates;
- **Malware Protection:** anti-malware solutions on systems accessing, processing, or storing Customer Data;
- **Audit Logs:** audit logs recording access to and actions taken upon Customer Data, sufficient to support security investigations and forensic analysis; and
- **AI-Specific Technical Controls:** technical measures specifically designed to prevent or mitigate AI Security Incidents, including: (i) input filtering and prompt injection defenses; (ii) output filtering to prevent leakage of Customer Data or Personal Data; (iii) controls preventing cross-customer contamination in multi-tenant model deployments; and (iv) controls limiting the risk of model inversion or extraction attacks.

Section 11.4: Physical Safeguards

Provider's physical safeguards shall include, at a minimum:

- **Facility Access Control:** limiting physical access to facilities where Customer Data is processed or stored to authorized personnel only;
- **Workstation Security:** controls to secure workstations and devices used to access Customer Data, including secure storage of portable devices and policies for remote work security; and
- **Media Security:** procedures for the secure handling, storage, transport, and disposal/destruction of physical media containing Customer Data to render it unrecoverable.

Section 11.5: Geographic and Time-Based Restrictions

Provider shall only access, process, and store Customer Data within [the United States] [the European Economic Area] [Specify Permitted Geographic Locations] and shall not transfer Customer Data outside of these locations without Customer's prior written consent. Provider may access Customer Data only during [specify times if applicable, e.g., business hours].

DRAFTER'S NOTE

Geographic Restrictions and Cross-Border Transfers. Geographic restrictions on Customer Data are increasingly important due to data localization laws (China, Russia, India) and cross-border transfer restrictions (EU adequacy decisions, GDPR Article 46). Counsel should: (a) identify the jurisdictions where Customer Data may be lawfully processed; (b) confirm Provider's ability to comply (some hyperscale cloud providers cannot meaningfully restrict processing locations); and (c) require notice and consent for any

change. For EU-origin data, counsel should also confirm that an appropriate transfer mechanism (e.g., Standard Contractual Clauses) is in place.

Section 11.6: Data Retention and Deletion

Provider shall retain Customer Data only for so long as, and only to the extent, necessary to provide the AI Services to Customer under this Agreement. Without limiting the foregoing:

(a) **Deletion Upon Instruction**. Upon Customer's written instruction at any time during the Term, and in any event upon the expiration or termination of this Agreement, Provider shall, within [thirty (30)] days, permanently delete or render permanently inaccessible all Customer Data in Provider's possession or control, including all copies, caches, logs, embeddings, vector representations, and other artifacts derived from Customer Data, except as expressly permitted under subsections (c) and (d) below.

(b) **Certification**. Upon Customer's request, Provider shall certify in writing, through an authorized officer, that Provider has completed the deletion required by subsection (a), identifying the categories of Customer Data deleted and the date of deletion.

(c) **Backup Media**. Where Customer Data resides on backup or archival media maintained in the ordinary course of Provider's disaster-recovery operations and cannot reasonably be deleted on an individual basis, Provider may retain such Customer Data until deletion occurs in the ordinary rotation of such media, not to exceed [ninety (90)] days following the deletion deadline in subsection (a), provided that during such period Provider shall: (i) not access, use, or restore such Customer Data for any purpose other than disaster recovery; and (ii) continue to protect such Customer Data in accordance with this Section 11.

(d) **Legal Retention**. Provider may retain Customer Data solely to the extent and for the duration required by applicable law, provided that Provider shall: (i) notify Customer of the legal requirement relied upon, to the extent permitted by law; (ii) isolate the retained Customer Data from active systems; (iii) not use it for any purpose other than compliance with the legal requirement; (iv) continue to protect it in accordance with this Section 11; and (v) delete it promptly upon expiration of the retention requirement.

(e) **Coordination**. Provider's obligations under this Section 11.6 are in addition to: (i) Provider's deletion and return obligations with respect to Customer-Specific Customizations under Section 4.2.3; and (ii) Provider's obligations to assist with data subject deletion requests under Section 12.3. For the avoidance of doubt, nothing in this Section 11.6 limits the prohibition in Section 2.1 on the incorporation of Customer Data into any AI Model; Provider's compliance with Section 2.1 is what ensures that deletion under this Section 11.6 is complete.

DRAFTER'S NOTE

Why Derived Artifacts Are Enumerated. Traditional deletion clauses address "copies" of data. In AI deployments, Customer Data also persists in forms that are not copies in the conventional sense: embeddings, vector stores, inference logs, and cached prompt/response pairs. Subsection (a) enumerates these expressly so Provider cannot take the position that deletion of source files satisfies the obligation while derived artifacts persist.

The Model Weights Question. This Section deliberately does not require deletion of Customer Data "from the AI Model," because (as noted in Section 2.3) removal of data integrated into model weights is generally technically infeasible. The architecture of this Toolkit handles the problem upstream: Section 2.1 prohibits Customer Data from entering

any training pipeline in the first place, so there should be nothing in the weights to delete. Subsection (e) makes this relationship explicit. Where Customer has authorized Customer-Specific Fine-Tuning under Section 2.2, the resulting fine-tuned artifacts are addressed by the destruction and certification obligations of Section 4.2.3, not this Section.

Backup Carve-Out Discipline. The backup exception in subsection (c) provides the vendor the most latitude; open-ended "backup rotation" language can extend retention indefinitely. The [90]-day outside limit, the no-restoration covenant, and the continued application of Section 11 safeguards are each load-bearing; counsel should resist deletion of any of them.

Regulatory Alignment. Counsel should confirm the periods nest properly for Customer's applicable regulatory regimes (e.g., GDPR).

Section 11.7: Security Event and AI Security Incident Notification

Provider shall notify Customer without undue delay, and in any event within [twenty-four (24)] hours of Provider's discovery, of: (a) any actual or reasonably suspected AI Security Incident affecting the AI Services or AI Model; and (b) any actual or reasonably suspected Security Event affecting Customer Data. Where an AI Security Incident also constitutes or reasonably may constitute a Security Event, Provider shall treat the notification obligation under clause (b) as triggered concurrently, and a single notification addressing both categories shall satisfy both obligations provided it contains all information required under each.

DRAFTER'S NOTE

Notification Timeline. The 24-hour notification window is a Customer-favorable starting position. Provider will typically push for longer (72 hours or “without undue delay” without a specific timeline). Customer should resist this because: (a) GDPR requires controller notification within 72 hours, which is impossible if Provider does not notify Customer promptly; (b) US state breach notification laws have varying timelines, some as short as 30 days; and (c) Customer’s reputational and operational response depends on rapid notification.

Distinct Treatment of AI Security Incidents. AI Security Incidents (defined separately in the Definitions section) may not constitute traditional Security Events because they do not involve unauthorized access to data in the conventional sense. However, they create distinct risks (cross-customer contamination, model integrity compromise) that warrant separate notification and response obligations. Treating them as a separate category ensures Provider cannot escape notification requirements by arguing that an AI-specific incident does not meet traditional breach definitions.

FALLBACK (PROVIDER-FAVORABLE NOTIFICATION TIMELINE)

Where Customer cannot achieve 24-hour notification:

“Provider shall notify Customer of any Security Event or AI Security Incident affecting Customer Data or the AI Services without undue delay, and in any event within [seventy-two (72)] hours of Provider’s confirmation of the event.”

Risk: Seventy-two hours may not provide Customer sufficient lead time to meet its own regulatory obligations, particularly under GDPR (72-hour controller notification window). Where this fallback is accepted, counsel should require Provider to provide preliminary notification immediately upon suspicion of an event, with confirmed details following within 72 hours.

Section 11.8: Audit Rights

Provider shall, upon reasonable advance notice and no more than [once per twelve-month period] (except where additional audits are required by Customer's regulators or are conducted following a Security Event or AI Security Incident), provide Customer or its independent auditor with: (a) access to information reasonably necessary to verify Provider's compliance with this Section 11; (b) access to Provider's relevant personnel for interviews; and (c) copies of Provider's most recent SOC 2 Type II reports, penetration test summaries, and similar security audit reports. Customer shall conduct any such audit during normal business hours and in a manner reasonably designed to minimize disruption to Provider's operations.

DRAFTER'S NOTE

Why Audit Rights Are Critical. Audit rights serve as the primary enforcement mechanism for the security and other obligations in this Agreement, particularly where contractual obligations are otherwise difficult to verify (e.g., the "insight laundering" concern under Section 4.3). Without audit rights, Customer's only remedy is breach of contract after the fact — which is often inadequate given the difficulty of detecting and proving security failures.

Provider Resistance. Providers strongly resist on-premises audit rights for many reasons (competitive concerns, third-party liability, operational disruption). Common compromises include: (a) reliance on SOC 2 Type II reports in lieu of on-site audits (acceptable for routine compliance but may not cover AI-specific risks); (b) shared assessor models where multiple customers participate in joint audits; and (c) third-party independent auditor models with NDAs. Customer should retain the right to additional audits in the event of a Security Event, AI Security Incident, or material change in Provider's practices.

Section 12: Data Privacy Issues to Consider

This Section addresses Provider's obligations regarding Personal Data and compliance with applicable privacy and data protection laws. AI fundamentally challenges traditional privacy assumptions, and this Section is structured to ensure legal defensibility and Customer control over how AI Services use and transform Personal Data.

DRAFTER'S NOTE

Why This Section Should Be Included. AI challenges traditional privacy assumptions in several important ways: (a) purpose limitation breaks down easily because AI systems are inherently designed to learn, generalize, and reuse data; (b) the data lifecycle is unclear because Personal Data may exist in training datasets, embeddings, logs, and outputs simultaneously; (c) controller/processor roles become blurred because some vendors act like processors contractually but behave like controllers operationally; (d) Outputs can create new regulated data by inferring sensitive traits, recreating Personal Data, or introducing bias or profiling risks; and (e) global regulatory scrutiny is increasing, with growing focus on automated decision-making, transparency, and data minimization.

Risks of Not Including This Section. Without robust privacy provisions, Customer faces: unauthorized use of Customer Data for AI training; regulatory liability shifts (Customer remains the "controller" under GDPR while Provider actions cause violations); inability to fulfill data subject rights (deletion, access, portability); indefinite or opaque data retention; hidden Subprocessor risks; AI-generated privacy violations through outputs that expose Personal Data or sensitive inferences; and PCI/regulated data exposure where applicable.

Vendor Negotiating Positions. Counsel should expect Provider to push back as follows: (a) broad "service improvement" rights to use Customer Data for model training; (b) de-identified/aggregated data carve-outs that risk re-identification; (c) "reasonable efforts" for data subject rights assistance instead of operational support; (d) retention flexibility for debugging, compliance, or backups; (e) Subprocessor opacity with minimal transparency into model providers or downstream processors; (f) silence on AI-specific processing (training, inference, output-related privacy); and (g) narrow definitions of "personal data" that exclude inferred or derived insights.

Section 12.1: Compliance with Data Protection Requirements

Provider shall comply with all applicable laws, regulations, and other legal requirements relating to privacy, data security, and protection of Personal Data (collectively, "Data Protection Requirements"). Data Protection Requirements include, but are not limited to: the EU General Data Protection Regulation (GDPR); the UK GDPR; the Swiss Federal Act on Data Protection (FADP); Brazil's Lei Geral de Proteção de Dados (LGPD); the California Consumer Privacy Act as amended by the California Privacy Rights Act (CCPA/CPRA); and other applicable U.S. state privacy laws.

DRAFTER'S NOTE

Other Jurisdictions to Consider. The Data Protection Requirements definition above is not exhaustive. Counsel should consider whether to add specific reference to: (a) Canada (PIPEDA / Bill C-27); (b) India (Digital Personal Data Protection Act 2023); (c) China (PIPL); (d) Singapore (PDPA); (e) Australia (Privacy Act); and (f) US federal sectoral laws (HIPAA, FERPA, GLBA, COPPA) where Customer operates in regulated industries.

Why Customer Might Want to Narrow the List. A broad, open-ended reference to applicable law or a long enumerated list creates uncertainty about which obligations apply and may bind Customer to compliance regimes it does not actually fall under. Provider will price and scope compliance based on what is in the Agreement, so a narrower list gives Provider less room to argue that unanticipated regulatory requirements are Customer's problem. Where Customer operates only in the U.S. with no EU data subjects, including GDPR may be unnecessary and may give Provider leverage to demand GDPR-level processing terms.

Why Customer Might Want to Keep It Broad. AI services often involve data flows that are difficult to predict at contracting time, so a broad definition future-proofs the Agreement. Customer may have global employees or customers whose data flows through the AI Services without Customer fully realizing it. Regulatory enforcement is also expanding rapidly, and narrowing the list today may create gaps tomorrow.

Section 12.2: Processing of Personal Data

If the Customer Data includes Personal Data, Provider shall process such Personal Data only as necessary to fulfill the purpose of this Agreement and consistent with data subject rights. Provider shall ensure that Personal Data is:

- processed lawfully, fairly, and in a transparent manner in relation to the data subjects;
- collected for specified, explicit, and legitimate purposes and not further processed in a manner that is incompatible with those purposes;
- adequate, relevant, and limited to what is necessary in relation to the purposes for which it is processed;
- accurate and, where necessary, kept up to date;
- kept in a form that permits identification of data subjects for no longer than is necessary for the purposes for which the Personal Data is processed; and
- processed in a manner that ensures appropriate security of the data, using appropriate technical and organizational measures (see Section 11).

Section 12.3: Data Subject Rights Assistance

Provider shall assist Customer in responding to data subject access, correction, deletion, portability, and other rights requests in compliance with applicable Data Protection Requirements. Such assistance shall include: (a) providing operational tools or processes to identify, retrieve, modify, or delete Personal Data of a specific data subject; (b) cooperating with Customer's investigations and responses to such requests within the timeframes required by applicable law; and (c) providing technical and operational support sufficient to enable Customer to meet its obligations as a data controller.

DRAFTER'S NOTE

Operationalizing Data Subject Rights. Many privacy laws require controllers to respond to data subject rights requests within tight timeframes (e.g., 30 days under GDPR). Where Provider acts as processor, Customer needs reliable, operational assistance to meet those timeframes. Vendor "reasonable efforts" language is often insufficient. Counsel should require: (a) specific tools or APIs to identify and extract Personal Data of a specific data subject; (b) a documented process and timeline for fulfilling requests; and (c) escalation paths for urgent requests.

AI-Specific Challenges. Data subject rights are particularly challenging in AI contexts because: (a) Personal Data may be embedded in model weights or embeddings, making true deletion technically infeasible; (b) Outputs may reproduce or infer Personal Data even after the underlying training data is deleted; and (c) right to explanation requirements (e.g., GDPR Article 22) may require Provider to provide information about model logic that Provider considers proprietary. Counsel should require Provider to address these challenges explicitly.

Section 12.4: PCI DSS Compliance

If Provider processes or transmits payment card information or Cardholder Data (as defined by the Payment Card Industry Data Security Standard ("PCI DSS")) on behalf of Customer, or is otherwise deemed in scope for PCI DSS compliance (including because Provider or Provider's information systems can impact the security of Customer's cardholder data environment, or because Provider is a "connected-to service provider" under PCI DSS), then Provider shall:

- comply with the most recent version of PCI DSS;
- upon Customer's request, provide Customer with an accurate, complete, and up-to-date Attestation of Compliance ("AOC") as defined by PCI DSS and compliant with payment card brand rules, covering the entire scope of Customer assets and Provider information systems that comprise Provider's cardholder data environment;
- acknowledge that Provider is responsible for the security of Cardholder Data that Provider possesses or otherwise stores, processes, or transmits on behalf of Customer, or to the extent that Provider could impact the security of Customer's cardholder data environment;
- comply with requests from Customer related to PCI DSS, including scoping exercises, targeted risk assessments, creation of a responsibility matrix, and fully cooperating with PCI DSS-required investigations for Security Events; and
- notify Customer if Provider experiences a significant change in its control environment as defined in PCI DSS.

If PCI DSS requirements are updated or changed by the issuing authority, Provider shall negotiate in good faith with Customer to amend this Agreement as necessary to allow Customer to comply with the amended PCI DSS requirements. In no event shall Provider allow its information systems or the AI Services to prevent Customer from obtaining and maintaining PCI DSS compliance.

DRAFTER'S NOTE

PCI DSS Scoping. PCI DSS scoping is technical and can be contentious. The "connected-to service provider" concept under PCI DSS v4.0 is intentionally broad and may apply even where Provider does not directly handle Cardholder Data. Counsel should engage Customer's PCI assessor (QSA) early in negotiation to confirm scope and obligations.

AI and Cardholder Data. AI Services trained on or used in connection with payment processing may create novel PCI compliance issues, particularly where the AI processes transaction data or fraud signals. Counsel should require Provider to confirm that AI Services within PCI scope do not introduce new exposure (e.g., Cardholder Data leakage through Outputs).

Exhibit A: AI Service Level Agreement

This Exhibit sets forth the service level commitments for the AI Services and the remedies for failure to meet those service levels. This Exhibit is referenced by and incorporated into Section 10 of the Agreement. Bracketed values are placeholders to be completed for each transaction.

DRAFTER'S NOTE

How to Use This Exhibit. This is a template, not a fixed schedule. The appropriate service level metrics, thresholds, and remedies depend heavily on the specific AI Services, Customer's use case, and the risk profile of the deployment. Counsel should: (a) select the metrics relevant to the use case; (b) negotiate specific quantitative thresholds based on Provider's capabilities and Customer's needs; (c) specify measurement methodology with sufficient precision to support enforcement; and (d) calibrate remedies (service credits, termination rights) to the materiality of the failure.

Provider Negotiation Dynamics. Providers typically resist quantitative AI performance commitments because: (a) AI accuracy is inherently probabilistic; (b) measurement methodology may be disputed; and (c) performance depends on Customer Inputs over which Provider has no control. Counsel should respond by: (a) accepting that commitments are statistical, not absolute; (b) specifying measurement methodology in detail; and (c) limiting commitments to designated task categories where measurement is feasible.

A.1: Availability

Provider shall maintain Availability of the AI Services of at least [99.9]% measured on a monthly basis. "Availability" is calculated as: (Total Monthly Minutes – Unscheduled Downtime Minutes) / Total Monthly Minutes × 100%. Unscheduled Downtime excludes scheduled maintenance windows announced at least [seven (7)] days in advance and emergency maintenance reasonably necessary to address security or stability issues.

DRAFTER'S NOTE

Availability Targets. Typical enterprise SaaS availability targets range from 99.5% (allowing roughly 3.6 hours of downtime per month) to 99.99% (allowing roughly 4.4 minutes per month). The choice depends on the criticality of the AI Services. For mission-critical use cases, consider 99.95% or higher; for non-critical use cases, 99.5% may be acceptable.

A.2: Latency

Provider shall maintain median end-to-end response latency for the AI Services of no more than [_____] milliseconds for [defined task category], and 95th percentile latency of no more than [_____] milliseconds, in each case measured at Provider's service ingress. Latency measurements exclude network transit time between Customer and Provider's service ingress.

DRAFTER'S NOTE

Latency Considerations. Latency commitments should be calibrated to: (a) the type of task (simple classification is much faster than long-form generation); (b) the deployment mode (synchronous user-facing applications need lower latency than batch processing); and (c) Customer's downstream system requirements. The median/p95 framework is standard but counsel may want to add p99 commitments for high-throughput applications.

A.3: Throughput

Provider shall support Customer throughput of at least [_____] requests per second sustained, and [_____] requests per second burst, for the AI Services. Provider shall provide [thirty (30)] days advance written notice of any reduction in supported throughput.

A.4: Accuracy Targets

For the designated task categories specified in Schedule A-1, the AI Services shall meet or exceed the following Accuracy Targets, measured according to the methodology specified in Schedule A-1:

- [Task Category 1]: minimum accuracy of [_____]% on the validation set specified in Schedule A-1;
- [Task Category 2]: minimum accuracy of [_____]% on the validation set specified in Schedule A-1; and
- [Additional task categories as applicable].

DRAFTER'S NOTE

Validation Sets. Accuracy commitments are meaningful only if the validation set is well-defined. Schedule A-1 should specify: (a) the source and composition of the validation set; (b) the labeling methodology; (c) the metrics used (accuracy, F1, precision/recall, BLEU, ROUGE, etc.); and (d) the frequency of re-validation. Where the validation set is sensitive (e.g., contains Customer Data), specify access controls and confidentiality protections.

A.5: Hallucination Rate

For task categories where applicable (as specified in Schedule A-1), the Hallucination Rate shall not exceed [_____]% measured according to the methodology specified in Schedule A-1. "Hallucination Rate" means the percentage of Outputs that contain factually inaccurate, fabricated, or unsupported claims, measured against the validation set.

A.6: Refusal Rate

For permitted task categories specified in Schedule A-1, the Refusal Rate shall not exceed [_____]% measured according to the methodology specified in Schedule A-1. "Refusal Rate" means the percentage of properly formed Inputs for permitted tasks that the AI Services decline or fail to process, excluding Inputs that violate the prohibited uses set forth in Section 6.2.

DRAFTER'S NOTE

Why Refusal Rate Matters. Overly conservative content filtering can render AI Services unusable for legitimate purposes. Refusal Rate commitments protect Customer from a situation where Provider tightens content filters to the point where the AI Services no longer perform their intended function. Counsel should also negotiate the right to escalate inappropriate refusals for review and adjustment.

A.7: Service Credits and Remedies

If Provider fails to meet any service level commitment in a given measurement period, Customer shall be entitled to service credits as follows:

- Availability below [99.9]% but above [99.5]%: [10]% of monthly fees attributable to the affected AI Services;
- Availability below [99.5]% but above [99.0]%: [25]% of monthly fees attributable to the affected AI Services;
- Availability below [99.0]%: [50]% of monthly fees attributable to the affected AI Services; and
- Failure of any other service level commitment for [two (2)] consecutive measurement periods: Customer may terminate the affected AI Services without penalty and receive a pro-rated refund of prepaid fees.

Service credits are not Customer's exclusive remedy. Customer may pursue any other rights and remedies available under this Agreement or at law for material or repeated failures to meet service level commitments.

DRAFTER'S NOTE

Non-Exclusive Remedy. The non-exclusivity language is critical. Providers typically argue that service credits are the "sole and exclusive remedy" for SLA breaches, which can leave Customer without recourse for material or systemic failures. Counsel should resist exclusive remedy language and ensure that Customer retains the right to terminate for material breach and pursue damages where service credits are inadequate.

A.8: Measurement Methodology and Reporting

Provider shall: (a) measure service level performance using methodology consistent with industry standards and the specifications in Schedule A-1; (b) provide Customer with monthly service level reports within [fifteen (15)] days of the end of each measurement period; and (c) provide Customer with reasonable access to underlying measurement data to permit Customer to verify service level performance. Disputes regarding service level performance shall be resolved by reference to Provider's underlying measurement data and the methodology specified in Schedule A-1.

Schedule A-1: Task Categories, Validation Sets, and Measurement Methodology

PLACEHOLDER — TO BE DRAFTED

This section (Schedule A-1) is a placeholder. Operative provisions, definitions, drafter's notes, and fallback positions to be drafted.

Subsections to address: Task category definitions specific to Customer's use case; Validation set specifications (source, composition, labeling methodology); Specific metrics for each task category (accuracy, F1, precision/recall, etc.); Measurement frequency and methodology; Hallucination measurement methodology; Refusal measurement methodology