ANTHROP\C

Claude Opus 4.8 Model Documentation Form

Model Documentation for Downstream Providers (EU AI Act Art. 53(1)(b), *Annex* XII; *GPAI Code of Practice*, *Transparency Measure* 1.1).

Date this document was last updated: July 24, 2026

Document version number: 2.0

General information

Legal name for the model provider:	Anthropic Ireland Limited
Model name:	Claude Opus 4.8
Release date:	May 28, 2026
Union market release date:	May 28, 2026
Model dependencies:	N/A

Model properties

Model architecture:	The model is a large-scale transformer model.

Input modalities:
For each selected modality please include maximum input size or write 'N/A' if not defined

X Text	**X Images**	**☐Audio**	**☐ Video**	**If any other please specify:**
Maximum size: 1 million tokens	Maximum size: 600 per request	Maximum size: N/A	Maximum size: N/A	Maximum size: N/A

Output modalities:
For each selected modality please include maximum output size or write 'N/A' if not defined

X Text	**☐ Images**	**☐Audio**	**☐Video**	**If any other please specify:**
Maximum size: 300,000 tokens	Maximum size: N/A	Maximum size: N/A	Maximum size: N/A	Maximum size: N/A

Methods of distribution and licenses

Method of Distribution	This model is available to downstream providers through the following distribution channels: 1. **Anthropic's Claude consumer applications:** Claude.ai web interface, desktop, and mobile applications (https://claude.ai). Subscription tiers include Free, Pro, and Max plans. 2. **Certain Claude APIs via Console:** (https://console.anthropic.com) and available marketplace listings and the Claude Agent SDK (https://code.claude.com/docs/en/agent-sdk/overview). 3. **Enterprise solutions:** Claude for Work (Team and Enterprise plans) via https://www.anthropic.com/enterprise. 4. **Third-Party Cloud Platform and Other Integrations:** Amazon Bedrock, Google Cloud Gemini Enterprise Agent Platform, and Microsoft Azure AI Foundry, in each case subject to such platforms' applicable terms in conjunction with Anthropic's applicable terms. 5. **Third Party Marketplaces:** Claude is available as an app or extension through marketplaces like Slack and Microsoft. 6. **Programming Interfaces and Integrated Development Environments:** Claude Code CLI tool and IDE integrations via extension https://docs.anthropic.com/en/docs/claude-code. All distribution methods provide black-box, inference-only access. Model weights are not made available.
Category of Licence	This model is provided under proprietary terms of service that restrict access to the model's source code and weights, and impose limitations on usage, distribution, and modification. Access for downstream use is provided through: ● Anthropic's Commercial Terms of Service (https://www.anthropic.com/legal/commercial-terms) or the negotiated enterprise

	agreement between Anthropic and the customer (in either case for developers using our API and Agent SDK, and for enterprise customers) • Applicable third-party platform terms (including for Amazon Bedrock, Google Vertex AI, and Microsoft Azure AI Foundry customers, in conjunction with Anthropic's applicable commercial terms).
Additional Assets	N/A

Use

Acceptable Use Policy:	Anthropic's Usage Policy: https://www.anthropic.com/legal/aup
Intended uses:	Anthropic provides Claude as a general-purpose AI assistant. We define uses that are restricted or prohibited in our Usage Policy (link above).
Type and nature of AI systems in	Anthropic provides Claude as a general-purpose AI assistant. We define uses that are restricted or prohibited in our Usage Policy (link above).

which the general-purpose AI model can be integrated:	
Technical means for model integration:	The model is integrated into AI systems via the Anthropic Messages API or through third-party cloud platforms (including Amazon Bedrock, Google Vertex AI, and Microsoft Azure AI Foundry). Required technical means: • API credentials: Obtained via Anthropic Console (https://console.anthropic.com) or partner platform • SDK integration: Official Python or TypeScript SDKs available via package managers (pip, npm) • Documentation: Comprehensive integration guides at https://docs.anthropic.com • Development tools: Anthropic Workbench for prompt development and testing Anthropic provides quickstart guides, code examples (Anthropic Cookbook), and reference implementations to assist integration. No specialised infrastructure beyond standard web connectivity is required for API access.
Required hardware:	N/A
Required sotware:	N/A

Information on the data used for training, testing, and validation

Data type/modality:
Select all that apply.

X Text	**X Images**	**☐Audio**	**☐ Video**	**If any other please specify:**

Data provenance:
Select all that apply

For definitions of each listed category, see the Template for the Public Summary of the Training Content of General-Purpose AI models provided by the AI Office

X Web crawling	X Private non-publicly available datasets obtained from third parties	X User data The model itself is not directly trained on user data, but auxiliary models such as preference models and classifiers may be. The model's training may incorporate data derived from Anthropic employees' use of internal-only model versions and from user-submitted feedback. The model may also be trained on synthetic data ultimately derived from user data.
X Publicly available datasets	X Synthetic data that is not publicly accessible (when created directly by or on behalf of the provider)	X Data collected through other means

		If any other please specify: A portion of the data corpus comes from acquired physical works, such as books.

Data curation methodologies:	All training data goes through both standard and proprietary internal processing. Generally, all data goes through a standard pipeline of de-duplication, encoding, and chunking. Lastly, we chunk the data into sequences of a fixed length for training the model (rather than training on a full document at a time). This can either decrease the length of a document (if longer than the sequence length) or put multiple documents in the same sequence (e.g., two short documents). We withhold some data that is used for evaluation over the course of training. We also manually look at random samples of the data to spot check for quality issues.