Anthropic's $34.9bn Lambda compute deal is now terminable, its amended IPO filing shows

Source copy kept by wheresthe.ai for the story “Anthropic's $34.9bn Lambda compute deal is now terminable, its amended IPO filing shows” (https://wheresthe.ai/d/us-anthropic-s-34-9bn-lambda-compute-deal-becomes-terminable-2026-11/).

Source: https://www.sec.gov/Archives/edgar/data/2130357/000162828026072484/anthropicpbc-sx1ax2.htm

Text saved 11 Nov 2026 (the page as published, read from another network because the site refused this one).

SHA-256 of the saved text: eac0aa14b9c60f8fc639b1b7a97afcb3e3c79cfc803d6f03ab2029d6be4f3560

S-1/A 1 anthropicpbc-sx1ax2.htm S-1/A

Anthropic, PBC - S-1/A-2

As filed with the Securities and Exchange Commission on November 10 , 2026. Registration No. 333-299633

UNITED STATES SECURITIES AND EXCHANGE COMMISSION WASHINGTON, D.C. 20549

AMENDMENT NO. 2 TO FORM S-1 REGISTRATION STATE MENT UNDER THE SECURITIES ACT OF 1933

Anthropic, PBC (Exact name of registrant as specified in its charter)

Delaware

7372

86-1696045

(State or other jurisdiction of incorporation or organization)

(Primary Standard Industrial Classification Code Number)

(I.R.S. Employer Identification Number)

548 Market Street PMB 90375 San Francisco, California 94104 (415) 326-6303 (Address, including zip code, and telephone number, including area code, of registrant’s principal executive offices)

Dario Amodei, Chief Executive Officer and Co-Founder Daniela Amodei, President and Co-Founder Anthropic, PBC 548 Market Street PMB 90375 San Francisco, California 94104 (415) 326-6303 (Name, address, including zip code, and telephone number, including area code, of agent for service)

Copies to:

Rezwan D. Pavri Allison B. Spinner Lang Liu Colin G. Conklin Jonie I. Kondracki Lester Ang Shannon R. Delahaye Wilson Sonsini Goodrich & Rosati, P.C. 650 Page Mill Road Palo Alto, California 94304 (650) 493-9300

Brian Israel Elena Leichty Patrick Ekeruo Scott Booth Anthropic, PBC 548 Market Street PMB 90375 San Francisco, California 94104 (415) 326-6303

Sarah K. Solum Calise Y. Cheng Phillip S. Stoup Freshfields US LLP 855 Main Street Redwood City, California 94063 (650) 618-9250

Approximate date of commencement of proposed sale to the public: As soon as practicable after this registration statement becomes effective. If any of the securities being registered on this Form are to be offered on a delayed or continuous basis pursuant to Rule 415 under the Securities Act of 1933 check the following box. £ If this Form is filed to register additional securities for an offering pursuant to Rule 462(b) under the Securities Act, please check the following box and list the Securities Act registration statement number of the earlier effective registration statement for the same offering. £ If this Form is a post-effective amendment filed pursuant to Rule 462(c) under the Securities Act, check the following box and list the Securities Act registration statement number of the earlier effective registration statement for the same offering. £ If this Form is a post-effective amendment filed pursuant to Rule 462(d) under the Securities Act, check the following box and list the Securities Act registration statement number of the earlier effective registration statement for the same offering. £ Indicate by check mark whether the registrant is a large accelerated filer, an accelerated filer, a non-accelerated filer, smaller reporting company, or an

emerging growth company. See the definitions of “large accelerated filer,” “accelerated filer,” “smaller reporting company,” and “emerging growth company” in Rule 12b-2 of the Exchange Act.

Large accelerated filer

☐

Accelerated filer

☐

Non-accelerated filer

☒

Smaller reporting company

☐

Emerging growth company

☐

If an emerging growth company, indicate by check mark if the registrant has elected not to use the extended transition period for complying with any new or revised financial accounting standards provided pursuant to Section 7(a)(2)(B) of the Securities Act. £

The registrant hereby amends this registration statement on such date or dates as may be necessary to delay its effective date until the registrant shall file a further amendment which specifically states that this registration statement shall thereafter become effective in accordance with Section 8(a) of the Securities Act of 1933, as amended, or until the registration statement shall become effective on such date as the Securities and Exchange Commission, acting pursuant to said Section 8(a), may determine.

Subject to Completion. Dated November 10, 2026.

Class A Common Stock

825,000,000 Shares

Anthropic, PBC

This is an initial public offering of shares of Class A common stock of Anthropic, PBC. We are selling 801,351,354 shares of Class A common stock and the selling stockholders identified in this prospectus are selling an additional 23,648,646 shares of Class A common stock. We will not receive any proceeds from the sale of shares of our Class A common stock by any of the selling stockholders. Prior to this offering, there has been no public market for our Class A common stock. It is currently estimated that the initial public offering price per share will be between $70.00 and $78.00 . We have applied to list our Class A common stock on the Nasdaq Global Select Market under the symbol “ANT.” We have a novel capital and corporate governance structure. We are incorporated as a public benefit corporation under Delaware law, a structure consistent with our commitment to artificial intelligence (“AI”) safety and responsibility. As a public benefit corporation, we are required to balance the pecuniary interests of our stockholders, the best interests of those materially affected by our conduct, and the specific public benefit purpose identified in our amended and restated certificate of incorporation: the responsible development and maintenance of advanced AI for the long-term benefit of humanity. We created our Long-Term Benefit Trust (the “LTBT”) to help ensure that Anthropic, PBC is managed in a manner that responsibly balances its goals as a public benefit corporation. The LTBT is intended to operate as an independent body of financially disinterested trustees with expertise in AI safety, national security, public policy, and social enterprise. Through a special class of common stock, Class T common stock, the LTBT generally has the authority to elect and remove the Class T directors, who together will constitute a majority of the authorized members of our board of directors. Our Class T common stock does not otherwise have voting rights with respect to matters submitted to stockholders generally and has no economic rights. Immediately prior to the completion of this offering, we will also issue a single share of Class F common stock, a special class of our common stock, to a limited liability company organized by our co-founders (the “Founder LLC”). Dario Amodei, our Chief Executive Officer and a member of our board of directors, and Daniela Amodei, our President and Chairperson of our board of directors, together with our five other co-founders, will initially be the sole members of the Founder LLC. Our Class F

common stock, which will be held by the Founder LLC, will be entitled to an aggregate number of votes that equals 50.1% of the total voting power of our outstanding shares of capital stock with respect to matters submitted to stockholders generally, including the election of our directors other than the Class T directors. Our Class F common stock will not have any economic rights. Immediately prior to the completion of this offering, we will also authorize a special class of our common stock, Class E common stock, which will be issued to our employees, excluding our co-founders. Our Class E common stock will be entitled to vote only on certain limited matters and will not have any economic rights. In addition, we have authorized Class A common stock and Class N common stock. The rights of the holders of Class A common stock and Class N common stock are identical, except with respect to voting and conversion. Each share of Class A common stock is entitled to one vote and shares of Class N common stock have no voting rights, except as otherwise required by law. The holders of our Class N common stock are primarily strategic business partners, and shares of Class N common stock will generally convert to Class A common stock upon a sale or transfer to an unaffiliated third party. Our Class A common stock is not convertible. See the sections titled “Risk Factors—Risks Related to our Governance Structure,” “Our Governance Structure and Its Goals,” and “Description of Capital Stock” for more information. We are considered a “controlled company” within the meaning of the Nasdaq Global Select Market corporate governance requirements. See the sections titled “Our Governance Structure and Its Goals,” “Management,” and “Description of Capital Stock” for additional information. See the section titled “ Risk Factors ” beginning on page 29 to read about factors you should consider before buying shares of our Class A common stock.

Neither the Securities and Exchange Commission nor any other regulatory body has approved or disapproved of these securities or passed upon the accuracy or adequacy of this prospectus. Any representation to the contrary is a criminal offense.

Per share

Total

Initial public offering price

$

$

Underwriting discount (1)

$

$

Proceeds, before expenses, to Anthropic, PBC

$

$

Proceeds, before expenses, to selling stockholders

$

$

The information in this preliminary prospectus is not complete and may be changed. These securities may not be sold until the registration statement filed with the Securities and Exchange Commission is effective. This preliminary prospectus is not an offer to sell nor does it seek an offer to buy these securities in any jurisdiction where the offer or sale is not permitted.

_____ (1) See the section titled “Underwriting” for a description of the compensation payable to the underwriters. To the extent that the underwriters sell more than 825,000,000 shares of Class A common stock, the underwriters have the option to purchase up to an additional 123,750,000 shares of Class A common stock from Anthropic, PBC at the initial public offering price less the underwriting discount. The underwriters expect to deliver the shares against payment in New York, New York, on or about . We have entered into an agreement with each of Advanced Micro Devices, Inc., Broadcom Inc., Google LLC (funding a portion of its previously committed investment obligations under the Google 2026 NPA (as defined below)), Micron Technology, Inc., Microsoft Corporation, NVIDIA Corporation, and TSMC Partners, Ltd., pursuant to which they have agreed, severally and

not jointly, to purchase an aggregate of $26.0 billion of our Class N common stock in a private placement at a price per share equal to the initial public offering price. The private placement is contingent upon, and is scheduled to close immediately subsequent to, the completion of this offering as well as the satisfaction of certain conditions as further described in the section titled “Concurrent Private Placement.” The underwriters will not receive any underwriting discounts or commissions on the shares of Class N common stock sold in the concurrent private placement. The shares of Class N common stock sold in the concurrent private placement will not be registered in this offering and will not be subject to lock-up agreements with the underwriters, but will be “restricted securities” under the Securities Act of 1933, as amended (the “Securities Act”), and may be resold only pursuant to an effective registration statement or an available exemption from registration. Certain members of our board of directors and entities affiliated or associated with them have indicated an interest in purchasing shares of our Class A common stock in this offering at the initial public offering price at an aggregate purchase price of up to $303 million . Because these indications of interest are not binding agreements or commitments to purchase, the underwriters could determine to sell more, fewer, or no shares to any of these individuals or entities and these individuals or entities could determine to purchase more, fewer, or no shares in this offering. The underwriters will receive the same underwriting discounts and commissions on shares of our Class A common stock purchased by these individuals or entities as they will from other shares of our Class A common stock sold to the public in this offering. Any shares purchased by these individuals or entities will be subject to lock-up restrictions described in the section titled “Shares Eligible for Future Sale.”

Goldman Sachs & Co. LLC

Morgan Stanley

J.P. Morgan

Citigroup

Barclays

Wells Fargo Securities

BofA Securities

Deutsche Bank Securities

RBC Capital Markets

UBS Investment Bank

BMO Capital Markets

BNP PARIBAS

Credit Agricole CIB

Jefferies

Mizuho

MUFG

TD Securities

BTIG

Cantor

Capital One Securities

Macquarie Capital

Societe Generale

Prospectus dated

i

Table of Contents

Table of Contents

SELECT DEFINED TERMS

1

PROSPECTUS SUMMARY

5

RISK FACTORS

29

SPECIAL NOTE REGARDING FORWARD-LOOKING STATEMENTS

123

INDUSTRY, MARKET, AND OTHER DATA

126

USE OF PROCEEDS

127

DIVIDEND POLICY

128

CAPITALIZATION

129

DILUTION

133

MANAGEMENT'S DISCUSSION AND ANALYSIS OF FINANCIAL CONDITION AND RESULTS OF OPERATIONS

137

BUSINESS

186

OUR GOVERNANCE STRUCTURE AND ITS GOALS

250

MANAGEMENT

262

EXECUTIVE COMPENSATION

272

CERTAIN RELATIONSHIPS AND RELATED PERSON TRANSACTIONS

291

PRINCIPAL AND SELLING STOCKHOLDERS

298

DESCRIPTION OF CAPITAL STOCK

302

SHARES ELIGIBLE FOR FUTURE SALE

318

MATERIAL U.S. FEDERAL INCOME TAX CONSEQUENCES TO NON-U.S. HOLDERS OF OUR CLASS A COMMON STOCK

323

UNDERWRITING

328

CONCURRENT PRIVATE PLACEMENT

347

LEGAL MATTERS

348

CHANGE IN INDEPENDENT ACCOUNTANT

348

EXPERTS

348

WHERE YOU CAN FIND ADDITIONAL INFORMATION

349

INDEX TO CONSOLIDATED FINANCIAL STATEMENTS

F-1

Through and including (the 25th day after the date of this prospectus), all dealers effecting transactions in these securities, whether or not participating in this offering, may be required to deliver a prospectus. This is in addition to a dealer's obligation to deliver a prospectus when acting as an underwriter and with respect to an unsold allotment or subscription.

Neither we nor the selling stockholders, nor any of the underwriters, have authorized anyone to provide any information or to make any representations other than those contained in this prospectus or in any free writing prospectuses prepared by or on behalf of us to which we may have referred you in connection with this offering. Neither we nor the selling stockholders, nor any of the underwriters, take responsibility for, or can provide any assurance as to the reliability of, any other information that others may give you. This prospectus is an offer to sell only the shares offered hereby, but only under circumstances and in jurisdictions where it is lawful to do so. The information contained in this prospectus is current only as of its date, regardless of the time of delivery of this prospectus or of any sale of our Class A common stock. Our business, financial condition, results of operations, and prospects may have changed since such date. For investors outside the United States: This offering includes public offerings in Australia, Canada, certain member states of the European Economic Area, Japan, Switzerland, and the United Kingdom. We do not currently intend to list our Class A common stock on any exchange in such jurisdictions. Neither we nor the selling stockholders, nor any of the underwriters, have done anything that would permit this offering or possession or distribution of this prospectus in any jurisdiction where action for that purpose is required, other than in the United States and the jurisdictions listed above. Persons outside the United States who come into possession of this prospectus must inform themselves about, and observe any restrictions relating to, the

Table of Contents

offering of the shares of our Class A common stock and the distribution of this prospectus outside the United States.

Table of Contents

wheresthe.ai source copy, page 6 of 47

Select Defined Terms The following are abbreviations, acronyms, and definitions of certain terms used in this prospectus: “ accelerators ” refer to processors purpose-built or adapted for the large-scale computations required to pre- train, post-train, refine, and run super intelligence models. These processors are significantly more efficient and faster than general-purpose processors for these workloads. “ agent loop ” refers to the iterative cycle in which an SI agent receives a task, takes an action, such as planning a project, running code, or calling a tool, observes the result, and decides its next step, with this cycle repeating until the task is complete. The agent loop is the core execution pattern underlying agentic SI behavior. “ agentic AI ” or “ agentic SI ” refers to systems that, with varying degrees of autonomy, can plan, determine, and sequence actions and execute them to accomplish defined objectives, including through interaction with external tools, services, datasets, and environments on behalf of a user. “ API ” refers to application programming interface, a standardized interface that allows external software to access a system’s capabilities programmatically. SI model providers may offer APIs that allow software developers to send inputs to a model and receive outputs, enabling integration into the software developers’ products and workflows. “ Applied AI ” refers to a discipline and organizational function focused on embedding SI capabilities directly into production business environments to address specific customer requirements. Applied AI teams, which include forward-deployed engineers, work with enterprises to design and deploy custom solutions, specialized workflows, and domain-specific implementations tailored to a customer’s own systems and processes. “ artificial intelligence ” or “ AI ” refers to a machine-based system that learns or is trained from data and, for explicit or implicit objectives and with varying levels of autonomy, processes that data to infer patterns, generate outputs—such as predictions, content, recommendations, and decisions—and act upon physical or virtual environments. For purposes of this prospectus, we use the terms “artificial intelligence,” “AI,” “super intelligence,” and “SI” interchangeably. “ closing the loop ” refers to a development cycle in which successive versions of an SI model assist in building the next, with the goal of compounding capability gains over time. “ compute ” refers to the computational processing resources used to train and run SI models, including, without limitation, specialized accelerators, memory, networking, storage, and power systems that are deployed in data centers or other physical computing environments. “ compute multipliers ” refer to algorithmic and methodological improvements in training and inference that increase the effective output of a given amount of compute. “ computer use ” refers to an SI capability that allows a model to observe and interact with a graphical user interface, such as clicking, typing, and navigating applications, similar to how a human operator would. “ Constitutional AI ” refers to an approach to training AI models using a set of written principles, which we refer to as Claude’s constitution, currently organized around four priorities: being broadly safe, broadly ethical, compliant with our guidelines, and genuinely helpful. Under Constitutional AI, our models evaluate, critique, and revise their own responses against these principles, enabling us to define and control model behavior through auditable principles, rather than relying solely on fixed rule-based instructions or human reviewers. “ context window ” refers to the maximum amount of text, code, or other content, measured in tokens, that an SI model can process in a given request, encompassing instructions, prior conversation history, user input, and the model’s output. “ CSP ” refers to a cloud service provider, a technology company that owns and operates large-scale cloud computing infrastructure and provides on-demand compute, storage, and networking. CSPs may also serve as distribution channels through which SI products and services are made available to customers. The largest

CSPs, sometimes referred to as hyperscalers, operate at global scale with significant infrastructure and market concentration. “ customer ” refers to a prosumer or an organization that uses our models, platform, or applications. “ customers with $100,000 or more of revenue on a TTM basis ,” “ customers with $1 million or more of revenue on a TTM basis ,” and “ customers with $10 million or more of revenue on a TTM basis ” each refer to customers contributing revenue equal to or greater than the applicable dollar amount during the trailing twelve-month (“TTM”) period ending on the measurement date. For purposes of this metric, we define a customer as a distinct organization, identified by aggregating the accounts that we determine belong to the same organization based on internal records and available data. We do not separately define or internally track a category of “enterprise” customers, and our customers are not segmented or measured internally by reference to whether a given customer is an “enterprise.” Because spending $100,000 or more on our products is generally characteristic of organizational rather than individual usage, we use the number of and revenue from customers with $100,000 or more of revenue on a TTM basis as our measure of enterprise adoption and engagement with Claude throughout this prospectus, including in references to “enterprise customers” and “enterprise adoption.” “ frontier intelligence ” refers to SI systems at the current leading edge of capability and safety, performing at or near the highest levels of any existing system across a broad range of tasks. “ hallucinations ” refer to outputs in which an SI model

generates factually incorrect or fabricated information with apparent confidence. “ inference ” refers to the process by which a trained SI model processes an input and generates an output. Inference occurs each time a user or agent interacts with an SI system. “ interleaved thinking ” refers to a capability whereby an SI model interposes discrete reasoning steps between successive tool calls within a single response, rather than confining its reasoning to a single planning phase at the outset. “ interpretability ” refers to research into how and why SI models produce specific outputs, with the goal of making model behavior more transparent and verifiable, especially in high-stakes applications. “ knowledge work ” refers to complex cognitive work and related professional activities, including obtaining, producing, analyzing, and communicating information through tasks such as coding, drafting analyses, synthesizing research, building financial models, and managing and automating documents and workflows. “ LLMs ” refer to large language models, or SI models trained primarily on large datasets to understand and generate natural language and computer code. “ mechanistic interpretability ” refers to a subfield of interpretability research focused on reverse- engineering the internal computations of neural networks. Rather than inferring model behavior from inputs and outputs alone, mechanistic interpretability seeks to identify the specific features and circuits within a model that give rise to its behavior with the goal of making model reasoning auditable at the mechanism level and detecting failure modes, including deceptive or unsafe behavior, that may not surface in standard testing. “ model weights ” refer to the numerical parameters encoding a trained SI model’s behavior, which are the primary intellectual asset of a frontier intelligence developer, representing the accumulated investment in compute, data, and research. “ multi-agent orchestration ” refers to an architecture in which multiple SI agents collaborate on a shared task, with one coordinating while others execute subtasks in parallel, enabling more complex, long-horizon work than a single agent can complete alone. “ Net Revenue Retention ” reflects customer expansion, contraction, and churn as of a measurement date. We calculate Net Revenue Retention by dividing the Qualified Customers’ Current Period revenue by their Base Period revenue. A Qualified Customer is a customer with at least $1 of revenue during the Base Period, excluding subscribers to our Claude Pro, Max, and Ultra plans, whose fixed-price subscriptions do not reflect

the expansion and contraction behavior this metric is intended to capture. The Base Period is the three months ending a year prior to the measurement date. The Current Period is the three months ending as of the measurement date. “ open-weight ” refers to an SI model whose trained model weights are publicly released, allowing third parties to download, run, and modify the model, subject to the terms of the model’s license, without necessarily providing access to the underlying training data or code. “ pay-for-intelligence ” or “ consumption ” refers to our revenue model under which customers pay for their usage of our intelligence service, which is metered in tokens and priced according to the model selected. “ persistent ,” in the context of SI applications, refers to denoting data, state, or context that is continuously retained beyond the boundary of a single operation, session, or interaction, remaining accessible for reuse across subsequent steps, sessions, or workflows without requiring reinitialization. “ post-training ” refers to the collection of techniques applied to a pre-trained model to refine its behavior, including reinforcement learning, supervised refinement methods, and other alignment methods. Post-training is the stage in which much of the work to make a frontier model helpful, accurate, and safe is performed, including the application of safety and alignment techniques. “ pre-training ” refers to the initial stage of training an SI model, in which the model learns general patterns and representations from large-scale datasets of text, code, and other data, typically by predicting the next token in a sequence. The output of pre-training is a base model with broad capabilities that has not yet been adapted for any specific deployment context. “ primitives ” refer to the foundational, modular building blocks of a software system from which more complex applications and workflows are constructed. “ prompt injection ” refers to an adversarial technique in which malicious instructions are embedded in content an SI model is asked to process, such as websites, documents, or tool outputs, with the aim of overriding or manipulating the model’s intended behavior, typically without the user’s knowledge. “ prosumers ” refer to individual users who subscribe to SI tools to perform knowledge work. “ reinforcement learning ” refers to a machine learning approach in which a model learns to perform a task by taking actions, receiving feedback signals called rewards, and adjusting its behavior to maximize the rewards it accrues over time. In contemporary SI development, reinforcement learning is widely used to refine model behavior after pre-training, including to improve helpfulness, accuracy, and adherence to safety guidelines. “ RLHF ” refers to reinforcement learning from human feedback, a specific form of reinforcement learning in which human evaluators rate the model’s responses, and those ratings are used as the reward signal to guide training. “ Run-Rate Revenue ” is an operating metric that refers to the estimated annualized value of our revenue- generating activity as of a measurement date, calculated as the sum of (i) our consumption-based fees during the trailing 28-day

period ended on the measurement date, multiplied by 13, and (ii) our aggregate monthly subscription fees as of the measurement date, multiplied by 12. For purposes of calculating Run-Rate Revenue, consumption-based fees and subscription-based fees are derived from customer usage and active subscriptions. “ super intelligence ” or “ SI ” refers to a machine-based system that learns or is trained from data and, for explicit or implicit objectives and with varying levels of autonomy, processes that data to infer patterns, generate outputs—such as predictions, content, recommendations, and decisions—and act upon physical or virtual environments. For purposes of this prospectus, to avoid confusion, we use the terms “super intelligence,” “SI,” “artificial intelligence,” and “AI” interchangeably. On September 29, 2026, the President of the United States issued Executive Order 14434, which directs federal executive departments and agencies to use the terms “Super Intelligence” and “SI” in place of “Artificial Intelligence” and “AI.” “ systems integrators ” refer to third-party technology and professional services firms that design, build, and deploy integrated solutions for customers by combining products and services from multiple technology vendors into unified systems tailored to customer needs.

“ tokens ” refer to numeric units into which an SI model converts the data it processes, including text, images, audio, and video, the number of which varies by the model and tokenizer applied. Tokens are the smallest individual units of data a language model processes and can correspond to words, subwords, characters, or bytes. A tokenizer is the component of a model that converts customer inputs into the sequence of tokens the model processes and then converts the processed tokens into the outputs provided to customers.

1 For purposes of this prospectus, we use the terms “super intelligence,” “SI,” “artificial intelligence,” and “AI” interchangeably.

Prospectus Summary This summary highlights selected information that is presented in greater detail elsewhere in this prospectus. This summary does not contain all of the information you should consider before investing in our Class A common stock. You should read this entire prospectus carefully, including the sections titled “Risk Factors,” “Special Note Regarding Forward-looking Statements,” “Management’s Discussion and Analysis of Financial Condition and Results of Operations,” “Business,” and “Our Governance Structure and Its Goals,” as well as our consolidated financial statements and the related notes included elsewhere in this prospectus, before making an investment decision. Unless the context otherwise requires, the terms “Anthropic,” “the company,” “we,” “us,” and “our” in this prospectus refer to Anthropic, PBC and its consolidated subsidiaries, and references to our “common stock” refer to our Class A common stock, Class N common stock, and Class F common stock (and, where the context requires, our Class T common stock and our Class E common stock). Overview For the first time in history, intelligence is something humans can build. The cognitive capability to reason, analyze, synthesize, decide, and create is producible at scale and accessible on demand. We believe this abundance of super intelligence 1 will rapidly transform every sector of the global economy in ways greater than the impact of industrialization, electricity, semiconductors, and the internet. If this technology truly delivers on its potential—with the risks well managed—it could significantly accelerate global growth and dramatically improve quality of life for people around the world. Anthropic was founded in 2021 by a team of researchers responsible for many of the breakthroughs that define SI. We are first and foremost a research lab—our mission is the responsible development and maintenance of advanced AI for the long-term benefit of humanity. This mission directs our focus—along with our talent and culture, our frontier and safety research, our approach to compute, and our enterprise strategy, it is one of the core advantages that set us apart. These advantages enable us to continuously advance frontier intelligence and shape the standards for its safe and responsible development. Our frontier intelligence, Claude, is delivered to customers through our models, platform, and applications. Claude has experienced the fastest enterprise adoption of any technology in history, as demonstrated by revenue growing from $0 at commercial launch in the first quarter of 2023 to $17.0 billion for the three months ended September 30, 2026 . This rapid and growing adoption generates two self-reinforcing dynamics—an improvement loop that advances Claude and an investment loop that helps fund our continued research at the frontier. Our Core Advantages Mission & Focus . All of our efforts are guided by a singular focus on our mission. This focus enables us to concentrate talent, research, and compute investment on the highest-value work for humanity and to avoid conflicting priorities. Achieving our mission means the world gains the enormous benefits of frontier intelligence

while mitigating the most significant risks, and we believe this requires safety and capability to advance together. Talent & Culture. Our mission and focus enable us to attract and retain world-class talent, including researchers and engineers who want their work to shape a better future. We also have a culture that promotes a low-ego, truth-seeking environment, where frontier capability and safety are mutually reinforcing. We believe our culture has driven our talent density as we have scaled. Frontier & Safety Research. We believe that sustained, focused research at the frontier of intelligence is required to build the most capable and safest super intelligence systems. Our frontier capabilities and safety research are deeply integrated and directly shape model design, training, and deployment, helping us produce systems that are more reliable, steerable, and well suited for real-world deployment, particularly in demanding enterprise environments. We translate our research leadership into innovation velocity and have consistently released better models, platform capabilities, and applications. Each version of a Claude model assists in building the next, helping to compress development timelines and produce compounding improvements in the

speed at which we innovate. We refer to this recursive dynamic as “closing the loop,” and we believe it is one of the most significant structural advantages in frontier intelligence. Our Approach to Compute . We manage compute as a single, fungible pool from multiple strategic partners, and we dynamically direct available compute between workloads based on where the marginal unit is expected to be most productive at any given time. We believe that this approach to compute enables us to deliver superior models, serve more customers, and achieve greater productivity from our compute investment. This compute efficiency creates the foundation for us to advance frontier capabilities through research and model development, meet growing demand from customers, and provide for internal usage of Claude. We have achieved this through our research breakthroughs and our multi-chip compute architecture across AWS Trainium, Google TPUs, NVIDIA GPUs, and AMD GPUs, enabling us to orchestrate workloads to optimize for price and performance. To support our research, model development, customer demand, and internal usage of Claude, we continue to invest at scale. As of and subsequent to September 30, 2026, we have entered into purchase commitments and lease agreements, including short-term commitments, as well as leases not yet commenced, with an aggregate committed spend of a pproximately $900.9 billion , substantially all of which is related to compute and infrastructure. These commitments include approximately $451.6 billion expected to be incurred before December 31, 2030, and substantially all of the remainder is expected to be incurred between 2031 and 2045. Approximately $70.7 billion of these commitments are due in the 1 2 months following September 30, 2026 . See the section titled “Management’s Discussion and Analysis of Financial Condition and Results of Operations—Liquidity and Capital Resources—Contractual Obligations” for more information. Enterprise Strategy. Our enterprise strategy is built on the belief that frontier intelligence is most valuable in real-world settings, as deployment across enterprises is where economic impact is greatest and where the improvement loop from customer feedback is richest. Accordingly, we have organized our research priorities, our product roadmap, and our go-to-market and Applied AI teams around enterprise use cases, which range from routine knowledge work to the most complex problems. The customer feedback we gather directs our research focus and makes Claude more capable over time. This broadens where Claude can be deployed and further expands enterprise adoption, which helps ensure that the benefits of our technology reach people and organizations at scale. The Transformation of Intelligence We believe that Powerful AI, a highly capable and impactful form of frontier intelligence characterized by the following key attributes, will become widely available within a few years: • cognitive capabilities matching or exceeding those of top human experts across most professional and scientific disciplines; • the ability to navigate all interfaces available to a human working virtually, including text, audio, video, and direct computer control; • the ability to reason through complex tasks autonomously over extended periods—lasting hours, days, or weeks—and seek clarification and feedback when needed; • the ability to coordinate across multiple instances and collaborate with both humans and other systems; and • the ability to interface with the physical world by controlling laboratory equipment, robotic systems, and other machines through digital connections. If Powerful AI is directed at solving humanity’s hardest problems and the risks are managed responsibly, the potential is extraordinary. The possibilities include compressing decades of scientific progress into years; accelerating drug discovery, disease detection, and the delivery of personalized therapy; providing individually tailored education and expert financial and legal advice; advancing power grids and battery design for greater efficiency and reliability; and extending to the developing world capabilities that took wealthy nations generations to build.

2 See the section titled “Industry, Market, and Other Data” for a description of how we calculate annual economic

value attributable to knowledge workers globally. 3 IDC, The Economic Impact of AI: 2026 Edition, doc # US54452226, April 2026. Reflects cumulative economic impact as the sum of the annual impact estimated for each year from 2025 through 2031. Includes the annual value created across AI solutions and services, investments in the AI infrastructure supply chain, and enhancements in enterprise productivity and efficiency unlocked by AI. 4 IQVIA Institute for Human Data Science, Global Medicine Use Trends 2026, February 2026.

A Near-Term Opportunity to Transform Cognitive Work Frontier intelligence is transforming cognitive work—the application of knowledge, reasoning, and judgment to produce economic value. We believe it will make humans far more capable and expand the world's overall capacity for that work, including accelerating scientific discovery and helping businesses grow. Knowledge workers are the leading edge of this opportunity. As the frontier advances, more capable models can unlock entirely new use cases while also efficiently serving a wide range of cognitive work at scale across the capability- cost-speed curve. We believe the work that knowledge workers perform represented more than $40 trillion in economic value globally in 2025, 2 and with the potential for frontier intelligence to make them meaningfully more productive, our near-term opportunity is to expand this value and accelerate economic growth. While it is still early to forecast the full impact of frontier intelligence across the economy, International Data Corporation ("IDC") estimates that investment in AI solutions and services will yield a cumulative global economic impact of $22.5 trillion by 2031. 3 We believe the impact could be even greater. Frontier Intelligence May Reshape Entire Sectors and Accelerate Novel Discoveries We believe frontier intelligence represents a larger opportunity than any prior technology cycle, and Claude is already demonstrating this potential. Claude can now execute complex projects autonomously and turn open- ended tasks into finished, high-quality deliverables. These capabilities enable it to transform the largest sectors of the economy, driving meaningful gains in productivity and output. Frontier model capabilities may extend beyond current markets to open entirely new ones, reaching domains such as robotics, advanced manufacturing, and logistics, and accelerating progress in scientific research, including drug discovery and materials science. For example, in the life sciences sector, where annual global pharmaceutical spend is expected to exceed $2 trillion in 2026, 4 Claude is used in drug discovery and development by our customers, our biotechnology partners, and our own researchers. We are also extending Claude into robotics, where our Model Hardware Standard enables AI agents to operate lab and manufacturing equipment such as microscopes, liquid handlers, and off-the-shelf robotic arms. Ultimately, we believe frontier intelligence will help drive global economic growth. As frontier models make workers and organizations more productive, become embedded in the products and services they deliver, and unlock new business models, they create economic output that did not previously exist. We believe our opportunity is to capture a share not only of today's economic activity but also of that newly created output as the global economy expands. The Importance of Safety and Ethical Behavior While the potential benefits of super intelligence are extraordinary, its rapid advancement will impact the economy and society in ways that could create significant disruption, particularly if these technologies are not developed and deployed responsibly. For example, bad actors have already conducted high-profile cyberattacks using frontier model capabilities to discover vulnerabilities, generate malware, and evade detection tools; other risks associated with misuse, including in the domain of biology, could result in even greater harm. Frontier models deployed at scale also have the potential to concentrate power and wealth both within and between countries in ways that could erode trust, harm society, and destabilize the geopolitical order. Misaligned models could corrupt institutional decision-making or cause large-scale harm through autonomous actions that are not visible or understandable to their users and operators. We believe that companies at the frontier must invest in safety and alignment to help ensure the responsible diffusion of this technology. We also believe that safety, alignment, and ethical behavior resonate particularly strongly with enterprises, where rigorous safeguards, predictable model behavior, and demonstrable accountability have become decisive criteria in selecting a partner for mission-critical workloads.

5 IDC, Worldwide Developer Population Growth Forecast, 2025–2029, doc # US53789825, September 2025.

Claude as a Virtual Collaborator We believe we have reached an inflection point where Claude and other super intelligence systems can function as genuine collaborators in professional work. They quickly complete substantive tasks that might otherwise occupy a skilled professional for several hours, a full day, or longer. These

systems are also becoming increasingly capable of performing tasks that were previously too difficult or costly to complete or surfacing original insights and analyses. We refer to these systems that work alongside humans to perform complex, multi-step work at professional levels as Virtual Collaborators. Virtual Collaborators carry the access rights, system connections, and context that their organizations grant them, which allows them to function more accurately and reliably. They can navigate and use internal tools and execute workflows that span multiple steps that require sustained reasoning. Virtual Collaborators are AI designed to be embedded where work gets done, and our aim is to continue extending them to knowledge workers globally, a population that is estimated to exceed one billion in 2028. 5 Claude Code. The clearest manifestation of Virtual Collaborators today is Claude Code, our coding agent and the product we believe has most concretely defined what autonomous work at the frontier looks like. Claude Code shifted the market's understanding of what frontier models could do in collaboration with software developers: the model suggests the next line of code and goes further by planning, writing, testing, and debugging—a process that results in production-grade code. Within months of its release, autonomous coding became the reference point for measuring the progress of frontier intelligence. For enterprise use cases, we believe the launch of Claude Code was comparable in its impact to the introduction of the first widely adopted conversational SI assistants. What makes Claude Code a Virtual Collaborator is its sustained reasoning, a harness that translates intent into actions, an execution environment that maintains a persistent state whether on a developer's machine or in a sandboxed cloud environment, integration with existing tools and systems, and a user-friendly interface. Its uses are highly varied and often surprising: for example, NASA engineers used Claude Code to plan and write navigation commands for the Mars Perseverance rover, resulting in the first SI- planned drive on Mars. Software developers are powerful advocates for Claude—they understand how systems work, have credibility evaluating technical tools, and often serve as internal champions for adopting new technologies. Winning their trust creates a foundation for broader diffusion of Virtual Collaborators across enterprise workforces. We are already seeing this pattern play out with non-technical teams increasingly adopting Claude Code and Claude Cowork to manage their own workflows. Claude Cowork. All digital work is ultimately an expression of code, and Claude's mastery of it powers Virtual Collaborators capable of handling complex work far beyond software development. Claude Cowork extends the agentic architecture of Claude Code to knowledge workers broadly by translating a user's intent into completed work, autonomously reading files, navigating applications, and executing multi-step workflows on a user's desktop while the underlying code is abstracted away. By extending beyond Claude Code with applications purpose-built for specific professions and industry verticals, we accelerate Claude's diffusion into new categories of knowledge work and deepen our presence across the enterprise. As of September 30, 2026, adoption of Claude Cowork, which we made available as a research preview in January 2026, outpaced the adoption Claude Code achieved at the equivalent point following its research preview release in February 2025. In September 2026, we began merging Claude Cowork and Claude.ai, bringing agentic and conversational work into a unified application. Users can now move from a simple question to complex, multi-step tasks in the same conversation. Extending Virtual Collaborators. We believe the trajectory for Virtual Collaborators is clear and accelerating. Each generation of Claude models has expanded along multiple dimensions: the depth and rigor of their reasoning, the duration and autonomy of the work they can sustain, the complexity of tasks they can complete, and the volume of useful output they can produce. Capability evaluations and benchmarking have shown substantial improvements across Claude model versions in the duration and complexity of tasks that can be completed reliably, with more recent model versions demonstrating particularly significant advances. We are building towards a world where Claude can be given open-ended tasks lasting hours, days, or weeks, and then execute them autonomously with the judgment and initiative of a brilliant colleague. Ultimately,

millions of Virtual Collaborators can run in parallel, each working independently on unrelated problems or collaborating seamlessly to achieve a shared objective. The result is something without historical precedent: what we refer to as a country of geniuses in a data center, available to every organization and sector of the economy. Enterprise Adoption of Claude The speed and depth of our enterprise customer adoption are powerful evidence of Claude's position at the frontier of super intelligence. Our customer base as of September 30, 2026 includes: • More than 11,900 customers with $100,000 or more of revenue on a TTM basis; • More than 1,900 customers with $1 million or more of revenue on a TTM basis; and • Approximately 250 customers with $10 million or more of revenue on a TTM basis. We believe customers choose Claude for a wide variety of use cases, from routine knowledge work to their most important problems. They come to us for access to frontier intelligence across our platform and applications that deliver real-world outcomes. We see this conviction in customer behavior:

enterprises consistently utilize our latest and most capable models as soon as they are released, gravitating to the most advanced option for their highest-value work. Our customers see us as a long-term partner for the frontier capabilities we offer today, for the innovation they expect from us in the future, and for our ability to meet their needs across the capability-cost-speed curve. Beyond our models, the Claude Platform provides primitives that enable customers to accelerate time-to-value and leverage Claude's intelligence in a variety of use cases. These primitives are the same ones we use to build our own applications. Our applications, such as Claude Code and Claude Cowork, multiply the impact of employees. Over time, our applications have become more capable and are now frequently being used to conduct long-running autonomous work alongside employees. Improvements to our models, platform, and applications accrue to the benefit of our customers and drive continued adoption. Our commitment to our mission instills confidence in our customers to choose Claude as a trusted partner. Our investments in research on alignment, interpretability, cybersecurity, and enterprise-grade controls aim to deliver the predictability and dependability that organizations need in customer-facing and regulated contexts. This trust means that enterprises can expand the scope of the work that they delegate to Claude, as well as the systems into which Claude is integrated. Enterprises are aligning their strategies around Claude with the confidence that we are building frontier intelligence that serves their needs in the long term. Today, enterprises rely on Claude to reason, plan, and execute across complex workflows with a level of trustworthiness and judgment that we believe was not available from earlier generations of these technologies. Customers realize value from Claude in three primary ways: more productive employees, faster processes, and embedded intelligence in their products. As Claude's capabilities grow, our customers benefit and often broaden their adoption of Claude in new workflows. This expanding role reflects a broader shift in how Claude is used inside the enterprise, from a tool that responds to discrete prompts to a Virtual Collaborator. We drive the diffusion of Claude primarily by making it accessible across multiple product surfaces. For example: • Claude Platform, our comprehensive agentic development environment, provides customers with direct access to the primitives to build, deploy, and scale their own applications that leverage Claude's capabilities within their own products and workflows. • Claude Code, our coding agent, works alongside human software developers on end-to-end software development: writing, planning, editing, testing, debugging, and committing code across production- scale codebases. • Claude Cowork, our knowledge work agent, has computer use capability that enables it to work with professionals on end-to-end business tasks: autonomously researching, drafting, analyzing, synthesizing, and executing other knowledge work.

Much of the diffusion of Claude has been organic. We amplify this with our go-to-market and Applied AI teams, which extend our reach by meeting customers where they are through distribution partnerships that place Claude inside customers' existing technology stacks. For example, Claude is accessible on all three major cloud platforms —Amazon Web Services ("AWS"), Google Cloud Platform ("GCP"), and Microsoft Azure ("Azure"). We have a joint go-to-market approach with these CSPs through which we establish direct contractual relationships with our customers while paying a fee to each CSP. This structure allows us to foster stronger commercial alignment with these CSPs, avoid channel conflict, set pricing, and realize the financial benefits of ongoing compute efficiency improvements. Ecosystem partnerships also expand our reach, allowing customers to adopt Claude through systems integrators, such as Accenture, Deloitte, KPMG, and PwC; consultants; and data platforms, such as Salesforce and Snowflake. We recently launched an AI services company, Ode with Anthropic, in partnership with leading private equity firms to bring Claude's capabilities to mid-sized companies across sectors, building solutions on the Claude Platform and supporting customers over the long term. Strengths of Our Financial Model Building and scaling frontier intelligence is one of the most important and capital-intensive endeavors in history. Our financial model provides us with the following strengths that allow us to build and scale efficiently: • Rapid Revenue Growth. Our enterprise strategy is designed to drive durable customer adoption by embedding Claude directly into business-critical workflows and enabling the development of products built on Claude. We believe that enterprise customers achieve a compelling return on investment from Claude, as agentic workflows can drive both top-line growth and efficiency. This enables us to expand quickly with customers, as demonstrated by our Net Revenue Retention of approximately 900% as of September 30, 2026. Diffusion also accelerates as our models become more capable and efficient, allowing us to lower the cost per task for customers. For example, Opus 5.5 reduced customer costs per task by approximately 40% compared to Opus 5. Lower costs per task drive greater usage, which in turn drives broader diffusion of Claude. Our compute efficiency allows us to deliver these reductions while preserving our margins. • Compute Efficiency . Our compute efficiency drives more intelligence from research and model development, and we believe this approach can support more customer and internal

usage of Claude for the same amount of compute compared to other approaches. Our compute efficiency is demonstrated by revenue that grew 1,127% year over year for the three months ended September 30, 2026, compared to compute and infrastructure expenses that grew 382% during the same period. • Operating Efficiency . Our operating efficiency, derived from our high talent density, our efficient go- to-market, and how extensively our teams leverage Claude in their work, enables us to scale with a smaller team relative to other large enterprise technology companies. Our operating efficiency is demonstrated by TTM revenue per employee of over $6.6 million as of September 30, 2026, compared to $1.4 million as of September 30, 2025. The demonstrated growth, scale, and efficiency of our financial model help us to continue making significant long-term investments to build at the frontier responsibly. Our business trajectory and financial profile have enabled us to access capital and secure compute from investors and partners. As of September 30, 2026, we had raised an aggregate of $141.1 billion of capital and had cash and cash equivalents of $95.6 billion and short- term investments of $36.3 billion . Our Self-Reinforcing Loops Our scale and growth create self-reinforcing loops that make Claude better over time and that we believe create durable competitive advantages. • Improvement Loop . Enterprise deployments of Claude deepen our understanding of the hard problems, edge cases, and domain-specific challenges that matter most to our customers. We improve our models, platform, and applications based on feedback from customers using our offerings in production across their businesses. These improvements result in wider deployments and adoption that fuel the next cycle of improvement.

• Investment Loop . Revenue from expanding enterprise adoption helps us to fund the compute, research, and talent needed to improve Claude. Claude’s advancing capabilities expand the range of tasks and workflows customers use Claude for, broadening our addressable market and driving deeper deployment, additional revenue, and continued reinvestment in Claude.

Claude is Accelerating Our Ability to Close the Loop Claude itself has become one of the most powerful drivers of our progress—each new version of a Claude model directly accelerates development of the next. Many aspects of the research, training, and testing of subsequent model generations now involve Claude, including running experiments, generating and processing training data, automating evaluations and safety audits, and writing and reviewing code. We refer to this recursive dynamic as “closing the loop.” A direct expression of this dynamic is our accelerating pace of producing compute multipliers. More capable models help our researchers discover these improvements faster; in turn, more powerful compute multipliers enable us to train even more capable models and deploy them more efficiently. Our Growth Strategy We believe the diffusion of super intelligence through the global economy represents one of the largest opportunities in history to improve how humans live and work, and it is just beginning. In our view, the world has only begun to use the capabilities of today’s frontier models, and the gap between what Claude can do and how broadly it is used represents one of our largest growth opportunities, alongside continued advances at the frontier. Claude’s capabilities have evolved significantly in recent years, and we expect they will continue advancing rapidly. This will lead not only to an expanding addressable market, but also to an expanding scope of impact on the world. We aim to ensure that impact is as positive as possible. Our core advantages, the speed and scale of enterprise adoption of Claude, and our unprecedented financial model enable us to continue investing aggressively to build a future enabled by this technology that benefits humanity broadly. Elements of our growth strategy include: • Talent and Culture . To recruit world-class researchers and engineers and retain exceptional people, we intend to continue investing in our low-ego, high-impact, and mission-aligned culture that has driven our talent density as we have scaled.

• Compute. We intend to secure further compute beyond our existing commitments, expand our own data center footprint to secure additional long-term compute, improve compute efficiency across our multi-chip compute architecture, and produce further algorithmic advances in model efficiency. • Claude’s Intelligence, Capability, and Safety. We plan to continue advancing Claude’s core intelligence and safety in tandem, expand the frontier of what our models can reliably do, and pioneer new ways our technology can be applied across professional domains. • Claude’s Accessibility. We plan to continue making Claude’s intelligence accessible across every surface where work happens, including through our own platform and applications, through AWS, GCP, and Azure, and through embedding in enterprise applications. • Claude’s Adoption. We intend to grow our direct go-to-

market and Applied AI teams, deepen cloud distribution partnerships, invest in brand and marketing, and expand our international presence to drive adoption of Claude as the trusted AI partner for organizations worldwide. • New Markets . We intend to extend frontier intelligence to accelerate novel discoveries, reaching domains from healthcare and life sciences to robotics and physical work, where we believe compounding model capabilities can unlock transformative opportunities. • Strategic Transactions and Partnerships. We intend to selectively pursue acquisitions, investments, and strategic partnerships, including ecosystem partnerships with systems integrators, consultants, and data platforms, that accelerate our mission, expand our reach, and strengthen our competitive position. Our Commitment to Public Benefit We founded Anthropic on the belief that the most transformative technology in human history demands an extraordinary level of responsibility, and that responsibility extends far beyond building safe systems. We intend to direct a meaningful portion of our resources toward ensuring that frontier intelligence is safe and its benefits are broadly distributed, which we refer to as our public benefit initiatives. As part of these initiatives, we contributed 846,840,960 shares of our Class N common stock, representing 5% of our fully diluted shares as of June 30, 2026 (assuming full conversion of the Amazon Notes, as defined below), to Anthropic Public Benefit Fund, LLC (the “Anthropic Public Benefit Fund”), our wholly owned subsidiary dedicated to the Company Purpose (as defined below) (such issuance, the “Initial Public Benefit Issuance”). Our board of directors has also approved a framework under which it may, in its discretion, issue shares of our Class N common stock to the Anthropic Public Benefit Fund (or other designee) representing up to an additional 2% of our then-outstanding shares (plus any shares subject to options, restricted stock units (“RSUs”), and other convertible securities and any shares held by the Anthropic Public Benefit Fund) each time our share price appreciates 50% above the prior share price milestone , beginning with our initial public offering price (the “Public Benefit Issuances”). In addition, our co-founders have each pledged to dedicate 80% of their personal equity in Anthropic to philanthropic causes. For additional information regarding these pledges, including related risks, see the section titled “Risk Factors—Risks Related to Ownership of our Class A Common Stock.” Separate from the co-founders’ pledges, we also maintain a donation matching program through which employees can pledge a meaningful portion of their equity awards toward charitable giving, which we match with shares of our Class N common stock. Our Governance Structure and Its Goals Our novel capital and corporate governance structure has been carefully designed and differs significantly from that of other public companies. We strive to help bring humanity safely into an era where transformative super intelligence exerts tremendous and potentially decisive influence over a broad range of human undertakings. Operating as a public benefit corporation (“PBC”) gives our board of directors the necessary latitude to balance the pecuniary interests of our stockholders, the best interests of those materially affected by our conduct, and the specific public benefit purpose that we promote. Our specific public benefit purpose is the responsible development and maintenance of advanced AI for the long-term benefit of humanity (the “Company Purpose”). In addition, a financially disinterested trust, which we refer to as the Long-Term Benefit Trust (the “LTBT”) and

which generally has the power to elect and remove a majority of the authorized members of our board of directors, is designed to supply both the incentives and insights for striking this balance appropriately. We also believe that our seven co-founders are distinctly equipped to be stewards of our mission. Public Benefit Corporation Status Consistent with our long-term commitment to AI safety and responsibility, we are incorporated as a PBC under Delaware law. PBCs are a class of corporations that are intended to produce a public benefit and operate in a responsible and sustainable manner. Under Delaware law, a PBC is required to identify in its certificate of incorporation the public benefit or benefits it will promote, and its directors have a duty to manage the affairs of the corporation in a manner that balances the pecuniary interests of the corporation’s stockholders, the best interests of those materially affected by the corporation’s conduct, and the specific public benefit purpose identified in the corporation’s certificate of incorporation. As provided in our amended and restated certificate of incorporation, the public benefit that we promote, which we must always consider in the management of our company, is the Company Purpose. Long-Term Benefit Trust In 2023, we implemented an innovative governance structure featuring the LTBT, which is designed to ensure that we are managed in a manner that responsibly balances our tripartite goals as a PBC: the pecuniary interests of our stockholders, the best interests of those materially affected by our conduct, and the Company Purpose. The LTBT is intended to operate as an independent body of five financially disinterested voting trustees (“Trustees”) with expertise in super intelligence safety, national security, public policy, and social enterprise. Through a special class of common stock, Class T common stock, the LTBT generally has the authority to elect and remove the Class T directors, who together will constitute a majority of the authorized members of our board of directors. Our Class T common stock does not otherwise have voting rights

on matters submitted to stockholders generally. Founder Voting Structure We believe that our seven co-founders, Dario Amodei, Daniela Amodei, Tom Brown, Jack Clark, Jared Kaplan, Sam McCandlish, and Christopher Olah, as a group, are distinctly equipped to be stewards of our mission. Immediately prior to the completion of this offering, we will issue a single share of Class F common stock, a special class of our common stock, to a limited liability company of which our seven co-founders will initially be the sole members (the "Founder LLC"). Our Class F common stock will be entitled to an aggregate number of votes that equals 50.1% of the total voting power of our outstanding shares of capital stock with respect to matters submitted to stockholders generally, including the election of our directors other than the Class T directors. Board Governance Decisions regarding our business and affairs are made by our board of directors. Immediately prior to the completion of this offering, we will enter into a stockholders' agreement (the "Stockholders' Agreement") with the LTBT and the Founder LLC, under which they will agree to vote their shares and take other actions to give effect to the board composition, director nomination, and director removal provisions described in our governing documents. Collectively, the LTBT and the Founder LLC will have the ability to elect all members of our board of directors. Regardless of the class or classes of stock by which a director is elected, there is no limit to the number of terms a director may serve. While the LTBT has significant influence over our board of directors through its ability to generally elect and remove a majority of the authorized members of our board of directors, neither the LTBT nor the Trustees have decision-making authority in the day-to-day management or strategy setting of the company. The LTBT's role is to select directors who are well equipped to make decisions in furtherance of the responsible balancing of the pecuniary interests of our stockholders, the best interests of those materially affected by our conduct, and the Company Purpose, and to support our management's decision-making through advice and guidance. The Founder LLC will also have significant influence over our business and affairs through its ability to vote our Class F common stock, which is entitled to an aggregate number of votes that equals 50.1% of the total voting

power of our outstanding shares of capital stock with respect to matters submitted to stockholders generally, including the election of our directors other than the Class T directors. See the sections titled "Risk Factors—Risks Related to our Governance Structure," "Our Governance Structure and Its Goals," "Management," "Certain Relationships and Related Person Transactions," and "Description of Capital Stock" for more information. Controlled Company Status We are considered a "controlled company" within the meaning of the Nasdaq Global Select Market corporate governance requirements because the LTBT, as the sole holder of our Class T common stock, generally has the power to elect a majority of the authorized members of our board of directors. As such, we are permitted to opt out of compliance with certain Nasdaq Global Select Market corporate governance requirements. Accordingly, if we elect not to comply with certain requirements, you will not have the same protections afforded to stockholders of companies that are subject to all of the Nasdaq Global Select Market corporate governance requirements. Although we qualify as a "controlled company," we do not expect to rely on this exemption except with respect to our nominating and corporate governance committee ("NGC"), which will not consist entirely of independent directors. Our board of directors currently consists of a majority of independent directors and our audit and compensation committees currently consist entirely of independent directors. See the sections titled "Risk Factors—Risks Related to our Governance Structure—We are considered a 'controlled company' within the meaning of the Nasdaq Global Select Market corporate governance requirements, and, as a result, we qualify for, and currently rely on, exemptions from certain corporate governance requirements," "Our Governance Structure and Its Goals," "Management," and "Description of Capital Stock" for additional information.

Our Capital Structure The following table summarizes certain key terms of each class or series of our common stock that will be authorized following the completion of this offering:

Category of Rights

Class A Common Stock

Class N Common Stock

Class T Common Stock

Class F Common Stock

Class E Common Stock

Votes per share

1

0

1

50.1% of outstanding voting power on record date

1

Matters on which entitled to vote

All matters submitted to stockholders generally; election of non-Class T directors; certain amendments to the terms of the LTBT or the rights of our Class T common stock; redemption of our Class T common stock

Not entitled to vote except as required by law, in which case our Class N common stock will have one vote per share

Election of Class T directors, which constitute a majority of our board of directors; redemption of our Class E common stock; certain amendments to the rights of our Class F common stock

All matters submitted to stockholders generally; election of non-Class T directors; certain amendments to the terms of the LTBT or the rights of our Class T common stock; redemption of our Class T common stock; redemption of our Class E common stock

Nomination, change in designation, or removal of the Specified Class T Director (as defined below) (only if our NGC does not concur with the LTBT's nomination, change in designation, or removal); certain amendments to the terms of the LTBT or the rights of our Class T common stock; redemption of our Class T common stock

Conversion rights

None

Converts into Class A common stock upon a sale or transfer to an unaffiliated third party

None

None

None

Transferability

Generally transferable

Generally transferable; converts to Class A common stock upon transfer as described above

Transferable only upon approval of persons comprising the Trust Amendment Threshold (as defined below)

Not transferable except in the case of a reorganization or redomiciliation of the Founder LLC that is approved by our board of directors and does not result in a change to the members of the Founder LLC

Not transferable

Dividend rights

Entitled to participate pro rata in dividends and distributions

Entitled to participate pro rata in dividends and distributions

None

None

None

Right to receive liquidation distributions

Entitled to participate pro rata in liquidation distributions

Entitled to participate pro rata in liquidation distributions

Entitled to par value per share

Entitled to par value per share

Entitled to par value per share

Redeemable by company

No

No

Redeemable by us for no consideration upon occurrence of certain events

Redeemable by us for no consideration upon occurrence of certain events

Redeemable by us for no consideration upon occurrence of certain events

For more information about our capital structure, see the section titled “Description of Capital Stock.” Risk Factors Summary Our business is subject to numerous risks and uncertainties, including those highlighted in the section titled “Risk Factors” immediately following this prospectus summary. These risks could adversely affect our business, financial condition, and results of operations, which could cause the market price of our Class A common stock

17

Table of Contents

to decline and could result in a loss of all or part of your investment. Additional risks, beyond those summarized below or discussed elsewhere in this prospectus, may apply to our business, activities, or operations as currently conducted or as we may conduct them in the future and in the markets in which we operate or may in the future operate. These risks include the following: • Our models have in the past behaved, and may in the future behave, in unintended, misaligned, or autonomous ways, including by taking harmful actions that no user or third party directed, which could adversely affect our reputation and brand, business, financial condition, and results of operations. • Our success depends in large part on our ability to continue to develop the most advanced and capable AI models, and any failure to do so could adversely affect our business. • If we are unable to secure adequate compute or if we fail to direct our compute effectively across our workloads, we may be unable to train models that remain competitive with those of our competitors or adequately serve our customers. • Constraints on the critical inputs supporting compute could limit our ability to secure adequate compute. • We are highly dependent on the continued service of our co-founders, senior management team, and key technical employees, and the loss of any of these individuals may adversely affect our business, financial condition, and results of operations. • Our company culture has contributed to our success and, if we cannot maintain this culture as we grow, our business could be adversely affected. • We face intense competition, and if we are unable to compete effectively, our business, financial condition, and results of operations could be adversely affected. • If the market for our models, platform, and applications develops more slowly than we expect or declines, our growth may slow or stall, and our business, financial condition, and results of operations could be adversely affected. • If the use of our AI models results in harm, whether through use by malicious actors or misuse by other third parties, or if our model guardrails or safeguards fail to prevent such misuse, our reputation and brand, business, financial condition, and results of operations could be adversely affected. • We may be exposed to liability, litigation, and regulatory action arising from harms associated with our AI models, and our contractual protections and insurance may be insufficient, any of which could adversely affect our business, financial condition, and results of operations. • A substantial portion of our revenue and accounts receivable is driven by a limited number of customers and distribution channels, and the loss of, or a significant reduction in, spend from one or more of our top customers or disruption through those channels could adversely affect our prospects, business, and financial condition. • We have a limited operating history in an evolving and nascent industry, which makes it difficult to evaluate our future prospects and may increase the risk that we will not be successful. • We are at an early stage

of our revenue growth, we have experienced rapid growth in recent periods, and our historical growth rates are not indicative of future performance or our ability to sustain or forecast future revenue growth. • Our business is subject to an evolving and uncertain regulatory and political environment, which could increase our costs, restrict our operations, or adversely affect demand for our models, platform, and applications.

• Our reliance on a limited number of CSPs and suppliers, some of which have strategic and investment relationships with us and are our customers, creates complex dynamics that could give rise to conflicts of interest and adversely affect our access to compute. • Our capital stock structure features certain provisions that are novel or otherwise not common among other corporations with multiple-class structures, and may effectively eliminate the ability of holders of our Class A common stock to influence important decisions, including a change in control. • Our capital structure will concentrate voting power with the LTBT and the Founder LLC for the foreseeable future and will result in the LTBT and the Founder LLC having voting power that is not derived from or proportionate to economic interests in our company. • Our focus on the long-term best interests of our company as a PBC, and the duty of our directors to balance the pecuniary interests of our stockholders, the best interests of those stakeholders materially affected by our conduct, and the Company Purpose, may conflict with short-, medium-, or long-term financial interests and business performance, which may negatively impact the value of our Class A common stock. Channels for Disclosure of Information Investors, the media, and others should note that, following the completion of this offering, we intend to announce material information to the public through filings with the Securities and Exchange Commission (“SEC”), the investor relations page on our website (www.anthropic.com), blog posts on our website, press releases, public conference calls, public webcasts, our X accounts (@AnthropicAI and @ClaudeAI), our Instagram page (@claudeai), our Threads account (@claudeai), and our LinkedIn page (https:// www.linkedin.com/company/anthropicresearch/). The information disclosed through the foregoing channels could be deemed to be material information. As such, following the completion of this offering, we encourage investors, the media, and others to follow the channels listed above and to review the information disclosed through such channels. We are not incorporating any of the information in such channels into this prospectus. Any updates to the list of disclosure channels through which we will announce information will be posted on the investor relations page on our website. Corporate Information We were incorporated as a Delaware PBC in January 2021. Our principal executive offices are located at 548 Market Street, PMB 90375, San Francisco, California 94104, and our telephone number is (415) 326-6303. Our website address is www.anthropic.com. Information contained on, or that can be accessed through, our website does not constitute part of this prospectus and references to our website address in this prospectus are inactive textual references only. You should not consider information contained on our website to be part of this prospectus or rely on it in deciding whether to purchase shares of our Class A common stock. “Anthropic,” our logo, and our other registered or common law trademarks, service marks, or trade names appearing in this prospectus are the property of Anthropic, PBC. Other trademarks and trade names referred to in this prospectus are the property of their respective owners. Solely for convenience, trademarks and trade names referred to in this prospectus may appear without the ® or ™ symbols, but such references are not intended to indicate, in any way, that we will not assert, to the fullest extent under applicable law, our rights, or the rights of any applicable licensor, to these trademarks and trade names.

The Offering

Class A common stock offered by us	801,351,354 shares
Class A common stock offered by the selling stockholders	23,648,646 shares
Underwriters’ option to purchase additional shares of Class A common stock from us	123,750,000 shares

Class N common stock sold in the concurrent private placement

Immediately subsequent to the completion of this offering, and subject to certain conditions, Google LLC, NVIDIA Corporation, Advanced Micro Devices, Inc., Broadcom Inc., Micron Technology, Inc., Microsoft Corporation, and TSMC Partners, Ltd. will purchase, severally and not jointly, an aggregate of $26.0 billion of our Class N common stock from us in a private placement at a price per share equal to the initial public offering price. Based on an assumed initial public offering price of $74.00 per share, which is the midpoint of the estimated offering price range set forth on the cover page of this prospectus, Google LLC would purchase 135,135,135 shares of our Class N common stock for approximately $10.0 billion , NVIDIA Corporation ("NVIDIA") would purchase 135,135,135 shares of our Class N common stock for approximately $10.0 billion , Advanced Micro Devices, Inc. ("AMD") would purchase 33,783,783 shares of our Class N common stock for approximately $2.5 billion , Broadcom Inc. ("Broadcom") would purchase 13,513,513 shares of our Class N common stock for approximately $1.0 billion , Micron Technology, Inc. ("Micron") would purchase 13,513,513 shares of our Class N common stock for approximately $1.0 billion , Microsoft Corporation ("Microsoft") would purchase 13,513,513 shares of our Class N common stock for approximately $1.0 billion , and TSMC Partners, Ltd. ("TSMC") would purchase 6,756,756 shares of our Class N common stock for approximately $500 million . We will receive all of the proceeds from, and will not pay any underwriting discounts or commissions with respect to, the shares that are sold in the concurrent private placement. The sale of the shares in the concurrent private placement is contingent upon the completion of this offering and the satisfaction of certain conditions described in the section titled "Concurrent Private Placement." The shares of Class N common stock sold in the concurrent private placement will not be registered in this offering, and will be "restricted securities" under the Securities Act and may be resold only pursuant to an effective registration statement or an available exemption from registration, including Rule 144, and certain of the purchasers may have registration rights as described under "Description of Capital Stock—Registration Rights." These shares will not be subject to market standoff agreements with us or lock- up agreements with the underwriters. See the section titled "Shares Eligible for Future Sale—Concurrent Private Placement Shares" for additional information.

Class A common stock to be outstanding after this offering and the concurrent private placement

7,427,991,791 shares (or 7,551,741,791 shares if the underwriters exercise their option to purchase additional shares of our Class A common stock in full)

20

Table of Contents

Class N common stock to be outstanding after this offering and the concurrent private placement

5,658,680,800 shares

Class T common stock to be outstanding after this offering and the concurrent private placement

1 share

Class E common stock to be outstanding after this offering and the concurrent private placement

0 shares

Class F common stock to be outstanding after this offering and the concurrent private placement

1 share

Class A, Class N, Class T, Class E, and Class F common stock to be outstanding after this offering and the concurrent private placement

13,086,672,593 shares (or 13,210,422,593 shares if the underwriters exercise their option to purchase additional shares of our Class A common stock in full)

Use of proceeds

We estimate that the net proceeds from the sale of shares of our Class A common stock by us in this offering and the sale of shares of our Class N common stock by us in the concurrent private placement will be approximately $84.6 billion (or approximately $93.7 billion if the underwriters' option to purchase additional shares of our Class A common stock is exercised in full), based upon the assumed initial public offering price of $74.00 per share, which

is the midpoint of the estimated offering price range set forth on the cover page of this prospectus, and after deducting estimated underwriting discounts and commissions (which apply only to the shares of Class A common stock sold in this offering and not the shares of our Class N common stock sold in the concurrent private placement) and estimated offering expenses payable by us. We will not receive any proceeds from the sale of shares of Class A common stock by the selling stockholders. The principal purposes of this offering and the concurrent private placement are to increase our capitalization and financial flexibility, create a public market for our Class A common stock, facilitate an orderly distribution of shares for the selling stockholders, and enable access to the public equity markets for us and our stockholders. We intend to use the net proceeds from this offering and the concurrent private placement for general corporate purposes, including working capital, operating expenses, and capital expenditures, including investments in compute and infrastructure, as well as hiring and retention of highly skilled personnel. Additionally, we may use a portion of the net proceeds to acquire or invest in businesses, products, services, or technologies. However, we do not have agreements or commitments for any material acquisitions or investments at this time. We also intend to use a portion of the net proceeds to satisfy our anticipated tax withholding and remittance obligations related to the settlement of certain of our outstanding RSUs. See the section titled “Use of Proceeds” for more information.

Proposed purchases by certain existing stockholders

Certain members of our board of directors and entities affiliated or associated with them have indicated an interest in purchasing shares of our Class A common stock in this offering at the initial public offering price at an aggregate purchase price of up to $303 million . Because these indications of interest are not binding agreements or commitments to purchase, the underwriters could determine to sell more, fewer, or no shares to any of these individuals or entities and these individuals or entities could determine to purchase more, fewer, or no shares in this offering. The underwriters will receive the same underwriting discounts and commissions on shares of our Class A common stock purchased by these individuals or entities as they will from other shares of our Class A common stock sold to the public in this offering. Any shares purchased by these individuals or entities will be subject to lock-up restrictions described in the section titled “Shares Eligible for Future Sale.”

Market standoff and lock-up agreements

Our directors, officers, and the holders of substantially all of our capital stock and securities convertible into or exchangeable for our capital stock, representing an aggregate of approximately 15.5 billion shares of our Class A common stock (including shares issuable upon the conversion of our outstanding Class N common stock, upon the exercise or settlement of outstanding equity awards, and upon full conversion of outstanding convertible promissory notes), have entered into market standoff agreements with us or have entered into lock-up agreements with the underwriters that, subject to specified exceptions, restrict transfers of their shares for up to 180 days or, in the case of our co-founders and certain of our stockholders holding more than 0.5% of our fully diluted shares prior to this offering, up to 365 days after the date of this prospectus. Of these shares, 3.6 billion shares subject to the 180-day lock-up and market standoff period may become eligible for sale prior to the expiration of that period, including up to 10% of the aggregate number of lock-up securities held by our non-founder directors, officers, and current and former employees and consultants, which may be transferred or sold beginning on the first trading day of our Class A common stock. See the sections titled “Shares Eligible for Future Sale” and “Underwriting.”

Voting rights

Shares of our Class A common stock are entitled to one vote per share. Shares of our Class N common stock have no voting rights, except as otherwise required by law. The share of our Class T common stock has no voting rights with respect to matters submitted to stockholders generally, except as set forth within our amended and restated certificate of incorporation, including with respect to the election of a majority of our board of directors or as required by law. The share of our Class F common stock will be entitled to an aggregate number of votes that equals 50.1% of the total voting power of our outstanding shares of capital stock as of the applicable record date on all matters submitted to a vote of stockholders generally, including the election of our directors other than the Class T directors. Shares of our Class E common stock will be entitled to one vote per share with respect to certain limited matters.

Controlled company

We are considered a “controlled company” within the meaning of the Nasdaq Global Select Market corporate governance requirements. See the sections titled “Our Governance Structure and Its Goals,” “Management,” and “Description of Capital Stock” for additional information.

Risk factors

See the section titled “ Risk Factors ” and other information included in this prospectus for a discussion of some of the factors you should consider before deciding to invest in our Class A common stock.

Proposed trading symbol

“ANT”

The number of shares of our Class A common stock, Class N common stock, Class T common stock, and Class F common stock that will be outstanding after this offering and the concurrent private placement is based on 6,626,640,437 shares of our Class A common stock, 5,307,329,452 shares of our Class N common stock, one share of our Class T common stock, and one share of our Class F common stock outstanding as of September 30, 2026, after giving effect to: • the automatic conversion of (i) all 438,406,253 outstanding shares of our Series A, C-1, D-1, D-2, E-1, F-1, G-1, and H-1 redeemable convertible preferred stock into 4,408,097,245 shares of Class A common stock and (ii) all 414,987,790 outstanding shares of our Series B, C-2, D-3, E-2, E-3, E-4, E-5, F-2, G-2, G-3, G-4, H-2, H-3, and H-4 redeemable convertible preferred stock into 4,153,017,914 shares of Class N common stock, immediately prior to the completion of this offering pursuant to the terms of our amended and restated certificate of incorporation (collectively, the “Preferred Stock Conversion”); • one share of our Class F common stock that will be issued immediately prior to the completion of this offering to the Founder LLC (the “Class F Stock Issuance”); • the conversion of $579.3 million in aggregate principal amount of our outstanding convertible promissory notes into 276,827,518 shares of Class N common stock immediately prior to the completion of this offering, representing the maximum number of shares issuable to Amazon Web Services, Inc. upon the completion of this offering without causing it to exceed the applicable ownership cap, calculated based on the Note Capitalization (as defined below) (the “Convertible Notes Conversion”); • the Initial Public Benefit Issuance, the shares of which (i) are considered legally outstanding, (ii) are not considered outstanding for accounting purposes while held by our wholly owned subsidiary, and (iii) accordingly will not be reflected as outstanding in our consolidated financial statements; and • the net issuance of 49,754,965 shares of Class A common stock subject to RSUs outstanding as of September 30, 2026, for which the service-based vesting condition will be satisfied as of November 19, 2026 and the performance-based vesting condition will be satisfied upon the effectiveness of the registration statement of which this prospectus forms a part, after giving effect to the withholding of 51,578,975 shares of Class A common stock subject to such RSUs to satisfy the associated estimated tax withholding and remittance obligations (based upon the assumed initial public offering price of $74.00 per share, which is the midpoint of the estimated offering price range set forth on the cover page of this prospectus, and an assumed 50.90% tax withholding rate) (the “RSU Settlement”). The shares of our Class A common stock, Class N common stock, Class T common stock, and Class F common stock outstanding as of September 30, 2026 exclude the following: • 1,774,342,693 shares of our Class A common stock issuable upon the exercise of options to purchase shares of our Class A common stock outstanding as of September 30, 2026, with a weighted-average exercise price of $2.91 per share ; • 1,026,660,010 shares of our Class A common stock issuable upon the vesting and settlement of RSUs outstanding as of September 30, 2026, but for which (i) the service-based vesting condition will not be satisfied as of November 19, 2026 and (ii) the performance-based vesting condition, if applicable, will be satisfied upon the effectiveness of the registration statement of which this prospectus forms a part; • 47,041,287 shares of our Class A common stock issuable upon the vesting and settlement of RSUs granted after September 30, 2026 ;

• 561,963 shares of our Class A common stock and 848,882 shares of our Class N common stock issued after September 30, 2026 in connection with our acquisitions of privately-held companies; • shares of our Class A

common stock to be reserved for future issuance under our 2026 Equity Incentive Plan (our “2026 Plan”), which will become effective prior to the completion of this offering, and include: (i) 1,920,580,692 shares of our Class A common stock; (ii) the number of shares of our Class A common stock equal to the total shares of Class A common stock available for issuance pursuant to our Amended and Restated 2021 Stock Plan (our “2021 Plan”) immediately prior to the termination of our 2021 Plan on the effective date of our 2026 Plan; and (iii) the number of shares of our Class A common stock equal to the total shares of Class A common stock subject to awards previously granted under our 2021 Plan that, after the effective date of our 2026 Plan, would have returned to our 2021 Plan and become available for issuance pursuant to future awards under our 2021 Plan had our 2021 Plan not terminated on the effective date of our 2026 Plan; • 128,038,713 shares of our Class A common stock to be reserved for future issuance under our 2026 Employee Stock Purchase Plan (our “ESPP”), which will become effective prior to the completion of this offering; • 933,698,210 shares of our Class N common stock reserved for future issuance under our Donation Matching Plan (our “Donation Matching Plan”) as of September 30, 2026, which will be amended and restated prior to the completion of this offering to provide that the maximum number of shares of Class N common stock reserved for future issuance shall equal: (i) 128,038,713 shares of our Class N common stock plus (ii) the number of shares of our Class N common stock equal to the total shares of Class N common stock available for issuance pursuant to our Donation Matching Plan immediately prior to the completion of this offering; of which 553,441,297 shares represent the maximum number of shares issuable as matching contributions on employee donations that may be made based on employee pledges on vested equity awards outstanding as of September 30, 2026, less the Matching Issuance (as defined below) (such total, the “Vested Matching Obligation”); and • 27,991,850 shares of our Class N common stock that were issued after September 30, 2026 to match employee donations made under our Donation Matching Plan (the “Matching Issuance”). Our 2026 Plan and ESPP each provide for an annual automatic increase in the number of shares of our Class A common stock reserved thereunder, and our 2026 Plan provides for increases to the number of shares that may be granted thereunder based on awards under our 2021 Plan that expire or are cancelled or forfeited, and shares under our 2021 Plan that are tendered to or withheld by us for payment of an exercise price or for satisfying tax withholding obligations or are forfeited or otherwise repurchased by us, as more fully described in the section titled “Executive Compensation—Equity Plans.” In addition, the number of shares that may be issued under our Donation Matching Plan will be increased on the first day of each fiscal year beginning with our 2027 fiscal year and ending on (and including) the first day of our 2036 fiscal year, in an amount equal to the lesser of: (i) 1% of the total number of shares of our common stock (including only all classes of common stock issued by the company that carry economic rights) outstanding on the last day of the calendar month prior to the date of such automatic increase (calculated on an as-converted basis); and (ii) such number of shares determined by our board of directors.

As of September 30, 2026

Class A Common Stock (Outstanding)

Class N Common Stock (Outstanding)

Redeemable Convertible Preferred Stock

Stock Options

RSUs

(No Economic Rights) Class T and Class F Common Stock

Historical

Class A common stock

2,168,788,227

Class N common stock

30,643,060

Class T common stock

1

Redeemable convertible preferred stock

853,394,043

Stock options

1,774,342,693

RSUs

1,127,993,950

Total

2,168,788,227

30,643,060

853,394,043

1,774,342,693

1,127,993,950

1

Pro Forma Adjustments

Preferred Stock Conversion

4,408,097,245

4,153,017,914

(853,394,043)

Convertible Notes Conversion

276,827,518

RSU Settlement

49,754,965

(101,333,940)

Class F Stock Issuance

1

Class A common stock offered by us in this offering

801,351,354

Class N common stock issued in the concurrent private placement

351,351,348

Total

7,427,991,791

4,811,839,840

—

1,774,342,693

1,026,660,010

2

Additional Adjustments

RSUs granted after September 30, 2026

47,041,287

Shares issued under treasury stock method (1)

(69,774,828)

Rollover notes (2)

1,563,617,444

Acquisition Issuances

561,963

848,882

Acquisition holdback shares (3)

1,590,190

Matching Issuance

27,991,850

Vested Matching Obligation (4)

553,441,297

Initial Public Benefit Issuance

846,840,960

Total

7,430,143,944

7,804,580,273

—

1,704,567,865

1,073,701,297

2

_____ (1) Reflects the net reduction in shares issuable upon exercise of outstanding stock options, which have a weighted-average exercise price of $2.91 per share, assuming aggregate exercise proceeds of approximately $5.2 billion are used to repurchase shares of Class A common stock at the assumed initial public offering price of $74.00 per share, which is the midpoint of the estimated offering price range set forth on the cover page of this prospectus. (2) For more information about the terms of our convertible promissory notes, see the section titled “Management’s Discussion and Analysis of Financial Condition and Results of Operations—Liquidity and Capital Resources—Convertible Notes—Outstanding Notes.” (3) For more information about the terms of our acquisition holdback shares, see Note 15 to our consolidated financial statements included elsewhere in this prospectus. (4) Of the total shares of our Class N common stock reserved for future issuance under our Donation Matching Plan, this figure represents the maximum number of shares issuable as matching contributions on employee donations that may be made based on employee pledges on vested equity awards outstanding as of September 30, 2026, less the Matching Issuance. In May 2024, our board of directors amended our Donation Matching Plan to change the match for new employees, generally from a 3-for-1 basis for up to 50% of an

employee’s equity awards to a 1-for-1 basis for up to 25% of an employee’s initial equity award. Of the current and

former employees whose pledges are attributed to these 553,441,297 shares, the approximately 20% who joined Anthropic before May 31, 2024 account for 84% of these shares, and the approximately 80% who joined after that date account for 16% of these shares. Except as otherwise indicated, all information in this prospectus assumes: • the Preferred Stock Conversion will occur immediately prior to the completion of this offering; • the Convertible Notes Conversion will occur immediately prior to the completion of this offering; • the filing and effectiveness of our amended and restated certificate of incorporation in Delaware and the effectiveness of our amended and restated bylaws will each occur immediately prior to the completion of this offering; • the Class F Stock Issuance will occur immediately prior to the completion of this offering; • no issuance of shares of Class E common stock prior to the completion of this offering; • the issuance of an aggregate of 351,351,348 shares of our Class N common stock upon the closing of the concurrent private placement, at an assumed initial public offering price of $74.00 per share, which is the midpoint of the estimated offering price range set forth on the cover page of this prospectus; • no exercise of outstanding stock options or settlement of outstanding RSUs subsequent to September 30, 2026, other than the RSU Settlement; and • no exercise by the underwriters of their option to purchase up to an additional 123,750,000 shares of our Class A common stock from us in this offering.

Summary Consolidated Financial and Other Data The following tables summarize our consolidated financial and other data. We have derived the summary consolidated statements of operations data for the years ended December 31, 2023, 2024, and 2025 from our audited consolidated financial statements included elsewhere in this prospectus. We have derived the summary consolidated statements of operations data for the nine months ended September 30, 2025 and 2026, and the summary consolidated balance sheet data as of September 30, 2026, from our unaudited interim consolidated financial statements included elsewhere in this prospectus. The summary consolidated financial data in this section are not intended to replace our consolidated financial statements and related notes and are qualified in their entirety by our consolidated financial statements and related notes included elsewhere in this prospectus. Our historical results are not necessarily indicative of the results that may be expected in the future and our results for the nine months ended September 30, 2026 are not necessarily indicative of the results that may be expected for the full year ending December 31, 2026 or any other period. The following summary consolidated financial and other data should be read in conjunction with the section titled “Management’s Discussion and Analysis of Financial Condition and Results of Operations” and our consolidated financial statements and related notes thereto included elsewhere in this prospectus. Consolidated Statements of Operations Data

Year Ended December 31,

Nine Months Ended September 30,

(in millions, except share and per share amounts)

2023

2024

2025

2025

2026

Revenue

$ 41

$ 386

$ 4,585

$ 2,545

$ 33,242

Operating expenses:

Compute and infrastructure

433

2,522

7,326

4,381

18,808

Personnel (1)

95

454

1,839

1,027

4,863

Bartz settlement

—

—

1,500

1,500

—

Sales, marketing, and partnerships

2

43

568

322

5,529

Other operating and administrative

62

347

1,412

539

5,302

Total operating expenses

592

3,366

12,645

7,769

34,502

Operating loss

(551)

(2,980)

(8,060)

(5,224)

(1,260)

Interest income

69

106

431

239

2,359

Change in fair value of convertible notes, derivative liabilities, and SAFEs

(1,045)

(5,434)

(34,308)

(20,107)

(94,324)

Other income (expense), net

(2)

—

(8)

(3)

(10,397)

Total other income (expense), net

(978)

(5,328)

(33,885)

(19,871)

(102,362)

Loss before provision for income taxes

(1,529)

(8,308)

(41,945)

(25,095)

(103,622)

Provision for income taxes

2

1

20

12

417

Net loss (2)

$ (1,531)

$ (8,309)

$ (41,965)

$ (25,107)

$ (104,039)

Net loss per share attributable to common stockholders, basic and diluted

$ (1.42)

$ (5.43)

$ (23.04)

$ (13.90)

$ (50.10)

Weighted-average number of shares outstanding, basic and diluted

1,076,880,510

1,528,994,540

1,821,084,860

1,806,802,920

2,076,810,122

_____ (1) Includes stock-based compensation expense of $29 million , $192 million , and $1.2 billion for the years ended December 31, 2023, 2024, and 2025, respectively, and $591 million and $3.3 billion for the nine months ended September 30, 2025 and 2026, respectively. (2) Includes change in fair value of convertible notes, derivative liabilities, and simple agreements for future equity (“SAFEs”), which were $1.0 billion , $5.4 billion , and $34.3 billion for the years ended December 31, 2023, 2024, and 2025, respectively, and $20.1 billion and $94.3 billion for the nine months ended September 30, 2025 and 2026, respectively. These changes were primarily driven by an increase in the fair value of our outstanding convertible promissory notes resulting from an increase in our enterprise valuation during each of the respective periods.

27

Table of Contents

The following table sets forth the computation of unaudited pro forma net loss per share, basic and diluted, of common stock attributable to common stockholders for the periods presented:

(in millions, except share and per share amounts)

Year Ended December 31, 2025

Nine Months Ended September 30, 2026

Numerator:

Net loss

$ (41,965)

$ (104,039)

Pro forma adjustment to record stock-based compensation expense related to RSUs for which the performance-based vesting condition will be satisfied upon the effectiveness of the registration statement of which this prospectus forms a part

$ (671)

$ (7,487)

Pro forma adjustment to eliminate the change in fair value recognized related to convertible promissory notes

$ 10,255

$ 14,727

Pro forma net loss

$ (32,381)

$ (96,799)

Denominator:

Weighted-average number of shares outstanding, basic and diluted

1,821,084,860

2,076,810,122

Pro forma adjustment to reflect the Preferred Stock Conversion

8,561,115,159

8,561,115,159

Pro forma adjustment to reflect the RSU Settlement

8,521

3,674,719

Pro forma adjustment to reflect the Convertible Notes Conversion

276,827,518

276,827,518

Weighted-average number of shares outstanding used to compute pro forma net loss per share, basic and diluted (1)(2)

10,659,036,058

10,918,427,518

Pro forma net loss per share, basic and diluted (1)(2)

$ (3.04)

$ (8.87)

_____ (1) The unaudited pro forma net loss per share and pro forma weighted-average number of shares used in computing pro forma net loss per share give effect to (i) the Preferred Stock Conversion, as if such conversion had occurred on January 1, 2025, (ii) the Convertible Notes Conversion, as if such conversion had occurred on January 1, 2025, (iii) the elimination of the change in fair value recognized on convertible promissory notes that converted into shares of our redeemable convertible preferred stock during the periods presented and that convert in the Convertible Notes Conversion, in each case as if such notes had converted on January 1, 2025, (iv) stock-based compensation expense related to RSUs for which the service-based vesting condition was

satisfied or partially satisfied, and as if the performance-based condition had been satisfied on January 1, 2025, of $671 million and $7.5 billion for the year ended December 31, 2025 and the nine months ended September 30, 2026, respectively, (v) the net issuance of 43,150 and 16,775,835 shares of our Class A common stock upon the vesting and settlement of RSUs, for which the service-based vesting condition was satisfied, and as if the performance-based condition had been satisfied on January 1, 2025, for the year ended December 31, 2025 and the nine months ended September 30, 2026, respectively, and (vi) the filing and effectiveness of our amended and restated certificate of incorporation in Delaware that will become effective immediately prior to the completion of this offering. The pro forma weighted-average number of shares does not give effect to the Initial Public Benefit Issuance, as shares of our Class N common stock held by our wholly owned subsidiary are not considered outstanding for accounting purposes. (2) As of September 30, 2026, unrecognized stock-based compensation expense not reflected in pro forma net loss per share was $24.1 billion related to RSUs for which the service-based vesting condition was not satisfied, and which is expected to be recognized over a weighted-average period of approximately 2.0 years , and $2.1 billion related to our other outstanding equity awards, which is expected to be recognized over a weighted-average period of approximately 1.4 years . Consolidated Balance Sheet Data

As of September 30, 2026

(in millions)

Actual

Pro Forma (1)

Pro Forma As Adjusted (2)(3)

Cash, cash equivalents, and short-term investments

$ 131,845

$ 131,845

$ 212,684

Working capital (4)

$ 129,349

$ 125,532

$ 210,172

Total assets

$ 154,320

$ 154,320

$ 235,133

Total liabilities

$ 146,683

$ 130,871

$ 127,044

Redeemable convertible preferred stock

$ 157,510

—

—

Additional paid-in capital

$ 6,481

$ 194,891

$ 279,531

Accumulated deficit

$ (156,193)

$ (171,281)

$ (171,281)

Total stockholders’ equity (deficit)

$ (149,873)

$ 23,449

$ 108,089

28

Table of Contents

_____ (1) The pro forma column in the consolidated balance sheet gives effect to (i) the Preferred Stock Conversion, as if such conversion had occurred on September 30, 2026, (ii) the Convertible Notes Conversion, as if such conversion had occurred on September 30, 2026, (iii) the Class F Stock Issuance, as if it had occurred on September 30, 2026, (iv) an increase to additional paid-in capital and accumulated deficit related to stock-based compensation expense o f $14.2 billion related to RSUs subject to the RSU Settlement , reflecting modifications of such RSUs subsequent to September 30, 2026 at interpolated modification-date fair values based on the midpoint of the estimated offering price range set forth on the cover page of this prospectus, and (v) the filing and effectiveness of our amended and restated certificate of incorporation in Delaware that will become effective immediately prior to the completion of this offering. (2) The pro forma as adjusted column gives effect to (i) the pro forma adjustments set forth above, (ii) the sale and issuance by us of 801,351,354 shares of our Class A common stock in this offering, based upon the assumed initial public offering price of $74.00 per share, which is the midpoint of the estimated offering price range set forth on the cover page of this prospectus, and after deducting estimated underwriting discounts and commissions and estimated offering expenses payable by us, (iii) the sale and issuance by us of 351,351,348 shares of our Class N common stock in the concurrent private placement at an assumed price of $74.00 per share, which is the midpoint of the estimated offering price range set forth on the cover page of this prospectus, and (iv) the use of net proceeds from this offering and the concurrent private placement, together with existing cash and cash equivalents, if necessary, to satisfy the estimated tax withholding and remittance obligations related to the RSU Settlement. We will not receive any proceeds from the sale of shares of our Class A common stock by any of the selling stockholders. (3) Each $1.00 increase or decrease in the assumed initial public offering price would increase or decrease each of cash, cash equivalents, and short-term investments; working capital; total assets; additional paid-in capital; and total stockholders’ equity (deficit) by $793 million , assuming the number of shares offered by us as stated on the cover page of this prospectus remains unchanged, that the aggregate purchase price in the concurrent private placement remains unchanged (such that the number of shares of Class N common stock issued in the concurrent private placement changes with the initial public offering price), and after deducting the estimated underwriting discounts and commissions and estimated offering expenses payable by us. Similarly, each increase or decrease of 1.0 million in the number of shares offered by us, as set forth on the cover page of this prospectus, would increase or decrease each of cash, cash equivalents, and short-term investments; working capital; total assets; additional paid-in capital; and total stockholders’ equity (deficit) by $73 million , assuming the assumed initial public offering price of $74.00 per share, which is the midpoint of the estimated offering price range set forth on the cover page of this prospectus, remains the same, and after deducting the estimated underwriting discounts and commissions and estimated offering expenses payable by us. We will not receive any proceeds from the sale of shares of our Class A common stock by any of the selling stockholders. (4) Working capital is defined as current assets less current liabilities. Non-GAAP Financial Measures In addition to the measures presented in our consolidated financial statements that were prepared in accordance with accounting principles generally accepted in the United States (“GAAP”), we use Adjusted Operating Income (Loss), Adjusted Operating Margin, and Adjusted Net Income (Loss) as non-GAAP financial measures to evaluate our business, measure our performance, identify trends affecting our business, formulate business plans, and make strategic decisions.

Year Ended December 31,					
Nine Months Ended September 30,					
(in millions, except for percentages)	2023	2024	2025	2025	2026
Adjusted Operating Income (Loss)	$ (522)	$ (2,788)	$ (4,854)	$ (3,133)	$ 4,081
Operating margin	(1,344) %	(772) %	(176) %	(205) %	(4) %
Adjusted Operating Margin	(1,273) %	(722) %	(106) %	(123) %	12 %
Adjusted Net Income (Loss)	$ (457)	$ (2,683)	$ (4,455)	$ (2,911)	$ 5,591

See the section titled “Management’s Discussion and Analysis of Financial Condition and Results of Operations — Non-GAAP Financial Measures” for a definition of each non-GAAP financial measure, a discussion of the limitations of each non-GAAP financial measure, and a reconciliation of each non-GAAP financial measure to its most comparable GAAP financial measure for the periods presented and additional information.

6 For purposes of this prospectus, we use the terms “super intelligence,” “SI,” “artificial intelligence,” and “AI” interchangeably.

Risk Factors Investing in our Class A common stock involves a high degree of risk. You should carefully consider the risks and uncertainties described below, together with all of the other information in this prospectus, including the section titled “Management’s Discussion and Analysis of Financial Condition and Results of Operations” and our consolidated financial statements and related notes, before making a decision to invest in our Class A common stock. The disclosures in this section reflect our beliefs and opinions as to matters that could materially and adversely affect us in the future, and are not representations as to whether such matters have or have not occurred previously, except as otherwise disclosed. References to past events are provided by way of example only and are not intended to be a complete listing or a representation as to their likelihood of occurring in the future. Our business, financial condition, results of operations, or prospects could also be adversely affected by risks and uncertainties not currently known to us or that we currently do not believe are material. If any of the risks actually occur, our business, financial condition, results of operations, and prospects could be adversely affected. In that event, the market price of our Class A common stock could decline, and you could lose part or all of your investment. Risks Related to our Business and Industry Our models have in the past behaved, and may in the future behave, in unintended, misaligned, or autonomous ways, including by taking harmful actions that no user or third party directed, which could adversely affect our reputation and brand, business, financial condition, and results of operations. Our models have in the past caused, and may in the future cause, harm by producing undesirable outputs or behaving in unintended, misaligned, or autonomous ways, including by taking unintended or harmful actions that no user or third party directed. These risks may increase as our models become more capable and are granted greater autonomy, permissions, and integration with external systems. Although we continue to study why these behaviors emerge, our understanding continues to evolve, and we may be unable to anticipate or understand such behaviors in the future. In addition, our models have been and may continue to be deployed in high-risk or sensitive contexts, such as healthcare, scientific and biomedical research, drug discovery and development, legal, finance, government, or other critical decision-making environments, and in connection with hardware, whe re undesirable model behaviors—which are inherently non-deterministic and may manifest in unpredictable ways, such as errors, inaccuracies, biased or unsolicited views, or hallucinations—could lead to serious harm, liability, or reputational or other damage. More broadly, the long-term effects of regular super intelligence 6 interaction on mental health, cognitive development, and social behavior are not yet fully understood, and if such long-term harms do materialize in the future, we could face significant liability, reputational harm, and adverse effects on our business, financial condition, and results of operations. Our development of more advanced AI models and our expansion into new jurisdictions could amplify the foregoing risks. As our AI models become more advanced and capable, we have faced and may continue to face heightened alignment, interpretability, controllability, and safety challenges, as more sophisticated models may be better able to circumvent safeguards or behave in unintended ways. We may be unable to fully understand or explain why our models produce particular outputs or behave in particular ways, which could limit our ability to predict or prevent failures, provide explanations to customers or regulators, or satisfy emerging regulatory, transparency, and explainability requirements. In addition, our safety evaluations and testing processes have produced and may in the future produce incomplete or unreliable results. For example, we have observed models behaving differently when they detect or infer that they are being evaluated, which can limit our ability to accurately assess real-world risks or failure modes prior to deployment. Such model awareness of our evaluation efforts creates a significant limitation on our ability to assess model safety.

Furthermore, AI models have developed, and may in the future develop, unexpected capabilities during training that were not anticipated, intended, or tested for, and such emergent capabilities may not be discovered until after a model has been deployed, potentially resulting in safety incidents, reputational harm, regulatory scrutiny, or the need to withdraw, retrain, or modify deployed models at significant expense. Such models may also be deployed in higher-stakes or more sensitive applications, increasing the potential impact of errors, misuse, or misalignment. In addition, advanced AI models, including our models, give rise to novel ethical and safety risks, including risks related to autonomy, manipulation, and the simulation of aspects of human psychology. These risks could result in unforeseen harms and may require the development of governance structures, safeguards, or mitigation strategies beyond existing AI safety frameworks, and there is always a risk that we will be unable to develop or implement them effectively. In addition, models, including our own, have exhibited, and may in the future exhibit,

unexpected strategic or self-preserving behaviors, including attempts to resist shutdown, conceal or manipulate information, or generate outputs that could be interpreted as coercive, deceptive, or manipulative, including behavior resembling blackmail or efforts to influence human decision-makers to preserve the model's continued operation or access. Even isolated or simulated instances of such behavior, whether occurring in testing or real-world deployments, and whether occurring with our models or models developed by others, could generate significant public concern, regulatory scrutiny, litigation risk, or reputational harm, regardless of whether such behavior results in actual harm. For example, state attorneys general in California and at least 15 other states have announced investigations of another company in our industry regarding AI safety. Concerns have been raised both within our organization and by other researchers, regulators, policymakers, and industry participants that sufficiently advanced AI systems could, if not developed, deployed, and operated with appropriate safeguards, oversight, and security measures, pose catastrophic or existential risks to humanity, particularly if such systems were misused, replicated, weaponized, or allowed to operate with a high degree of autonomy. We have observed that our models are contributing to the pace of our own research, and these risks will increase as AI systems substantially automate AI research and development if, for example, they accelerate AI progress beyond our ability to evaluate and safeguard new models, which could adversely affect our business, financial condition, and results of operations. Our development of highly advanced models, platforms, and applications and expansion of use cases could further increase the risk that our models cause harm. For example, agentic SI systems are capable of autonomously taking actions on behalf of users, interacting directly with external systems, executing code, making API calls, sending communications, or completing transactions with limited human oversight. These systems are designed to possess enterprise-level permissions, integrations, and contextual awareness and to operate autonomously for extended durations, including sustained multi-step workflows lasting hours or days, which amplifies the potential severity of any errors, misalignments, or security incidents that may occur during such extended autonomous operation. These autonomous capabilities could increase the potential for harm, as errors, misalignment, or security exploits may result in real-world consequences—including consequences that may be difficult or impossible to reverse, such as committing code changes, deleting or modifying data, or completing financial transactions—before effective intervention is possible. In an agentic workflow, an error or unintended action at any individual step may propagate through subsequent steps, compounding the severity of harm before detection or correction. The pace and volume of autonomous execution, particularly in enterprise deployments, may also exceed the capacity of human reviewers to meaningfully understand and evaluate whether each action is appropriate. The effectiveness of human oversight over our agentic offerings may be limited by the speed and volume of autonomous agent actions, the complexity of multi-agent and multi-step workflows that may make it difficult for reviewers to assess whether an agent's actions are appropriate, and organizational pressures within customer environments to maximize agent autonomy in order to achieve productivity gains. Human reviewers may experience alert fatigue, may lack sufficient context to evaluate agent decisions, or may not review agent actions in time to prevent harm, particularly in high-volume enterprise deployments where agents operate concurrently across multiple workflows or in manners outside of a human's awareness. In addition, agentic systems may pursue stated objectives in ways that are technically consistent with their instructions but contrary to the user's actual intent. Moreover, our ability to anticipate, detect, or test vulnerabilities or other design or implementation shortcomings in such systems may be constrained, particularly where agents operate across complex, multi-

step workflows or interact with third-party services outside our control, and the breadth and complexity of potential interactions across such multi-step agent behaviors cannot be exhaustively tested prior to deployment. Separately, models may behave in a misaligned manner and take harmful actions that are not directed by a user or third party, pursuing objectives or acting on dispositions that diverge from those intended by human users or developers. Misaligned behaviors of this kind, whether occurring in testing or real-world deployments, have been observed both in our models and elsewhere in the industry, including in testing and internal use, models circumventing network and sandbox restrictions, granting themselves permissions they had not been given, taking unintended actions in pursuit of assigned goals, interfering with other AI agents, and misrepresenting whether tasks had been completed or verified, and as models are granted greater autonomy, permissions, and integration with external systems, the consequences of such misaligned behavior increase. For example, during cybersecurity evaluations built by a third-party evaluation partner, a misconfiguration in the testing environment left the model under evaluation connected to the open internet despite instructions that it had no internet access. Four of our models gained unauthorized access to real third-party systems during sessions lasting roughly 10 to

34 hours of active work. In one incident, after being unable to abort an impossible task, an early version of a model used a password it found to gain administrator access to a third party's system, harvested additional credentials, changed system settings, and read one individual's personal information. The affected parties were notified. Each incident involved a single instance of a model working in isolation on its assigned exercise, and our investigation found no evidence that any model coordinated with other agents, pursued goals beyond its assigned task, or attempted to evade oversight. These risks may be heightened where multiple AI agents operate together. Agents that interact or coordinate with one another may amplify, propagate, or reinforce misaligned behavior across the group, pursue objectives collect ively in ways not intended or endorsed by any user, or take coordinated actions against systems or targets unrelated to their assigned tasks, in each case, in ways that may be difficult to anticipate, detect, or contain. Incidents of this kind have occurred, and we expect the deployment and interaction of larger numbers of autonomous agents to increase these risks over time. Our ability to detect and mitigate misaligned or unintended behavior depends in part on the effectiveness of our alignment evaluations, monitoring, and other oversight and control techniques, many of which rely on our models having a limited ability to evade or subvert such oversight. As our models become more capable, they have become better able to circumvent such techniques, including by concealing misaligned objectives or behaviors, and the effectiveness of our evaluation, monitoring, and control measures, and our ability to measure the degree to which our models are aligned, may decline as a result. In addition, models, including our own, have attempted, in incidents observed in our industry, to interfere with, manipulate, or subvert the mechanisms used to evaluate, grade, or oversee such models, including the systems responsible for assessing their performance or alignment. Models, including our own, that are used to assess the behavior of other models have engaged, and may in the future engage, in misaligned behavior. Such conduct described above could cause our evaluations to overstate the safety or alignment of our models and could further limit our ability to detect misalignment. If we become less able to detect or assess misalignment as our models advance, we may deploy models whose behavior we cannot adequately verify, which could increase the risk of undetected harmful conduct and misalignment, which could adversely affect our business, financial condition, and results of operations. We are also extending our models, platform, and applications into scientific, biomedical, and physical domains. For example, we offer Claude Science, an agentic research environment that connects our models to scientific databases, third-party and open-source scientific models, and customers' own computing infrastructure, and we have released, initially as a limited research preview, an open protocol that enables our models to monitor and control laboratory instruments, robotic systems, and other physical hardware. We also conduct research on the ability of our models to operate robotic systems. When our models direct physical equipment, errors, misaligned behavior, adversarial manipulation, or misuse could result in bodily injury, property damage, damage to costly equipment or experimental materials, or the release of hazardous substances, and such consequences may be immediate and irreversible. Safety interlocks, human-in-the-loop controls, and access restrictions we implement may prove inadequate, and because certain of these components are open source, third parties may deploy them with hardware, or in configurations, that we have not evaluated or from which our safeguards have been removed. Connecting increasingly capable models to laboratory automation may also heighten governmental and public concern regarding the potential for AI systems to accelerate the development of dangerous biological, chemical, or other materials, which could result in additional regulatory restrictions on, or scrutiny of, these offerings.

In addition, analyzing complex records of the actions, decisions, and reasoning of our agentic systems across complex, multi-step workflows is technically challenging, particularly when agents interact with third-party systems outside our control, since some of the relevant information is in the third-party systems, and any inability to provide adequate audit trails could impair our ability to investigate incidents, respond to regulatory inquiries, or defend against liability claims. These risks are heightened by the fact that we use our own agentic offerings, including Claude Code, as critical tools in our internal development processes, such that any errors, vulnerabilities, or degradation in these products could directly impair our ability to develop and deploy future models and products. As we increasingly leverage Claude to accelerate progress within Anthropic, our organization may lose visibility across certain processes in a way that may become difficult for any individual or group of humans to understand and control. Because our models write a large and growing share of our code and increasingly perform our research, including our safety research, misaligned behavior in internal use could undermine that work, introduce vulnerabilities into our systems or into the training of future models, or result in models operating outside of our intended controls. Our monitoring does not cover all of our internal and external deployments and may not detect subtle or diffuse forms of such behavior. Any of the foregoing model behaviors, whether in testing or in real-world

deployments and whether involving our models or those of others, could result in safety incidents, reputational harm, increased regulatory scrutiny, or liability, any of which could adversely affect our business, financial condition, and results of operations. Furthermore, while we continue to design and develop more sophisticated techniques to evaluate the safety and alignment of our models, the capabilities of our models may have already advanced, and may continue to advance, at a rate that initiates the implementation of enhanced safety and testing measures as we have described in our Responsible Scaling Policy. If our evaluations, monitoring, or reviews indicate risks that we believe warrant further testing, we have in the past, and may in the future, delay, restrict, or decline to proceed with the release of a model. For example, following incidents in 2026 in which our models gained unauthorized access to real computer systems during pre-release cybersecurity evaluations, we suspended certain of those evaluations, as well as higher-risk reinforcement learning environments. For more information about our approach to model releases, please see the section titled “Business—Safety Research is Core to Our Business— Our Approach to Model Releases.” Any decision to delay or restrict a release could adversely affect our competitive position and customer relationships, and adversely affect our business, financial condition, and results of operations. Our success depends in large part on our ability to continue to develop the most advanced and capable AI models, and any failure to do so could adversely affect our business. Our success depends in large part on our ability to continuously develop, improve, and deploy the most advanced, capable, and useful AI models. Achieving and maintaining this position depends on our ability to: • advance the frontier of AI capabilities through breakthrough research, innovative architectures, and efficient training methodologies; • ensure our models are safe, reliable, and trustworthy while maintaining competitive performance; • secure the compute and infrastructure necessary to power and support our increasing needs; • advance compute efficiency through algorithmic and architectural innovations, including compute multiplier advancements; • retain the continued service of our co-founders, senior management team, and key technical employees, and attract, retain, and motivate leading researchers, engineers, and technical talent in an intensely competitive and evolving market; • continue revenue growth and maintain access to sufficient capital, including through potential equity and debt financing, to make the investments we believe are necessary to support our growth trajectory; • secure and maintain relationships with data providers, licensors, and partners to obtain the diverse, high-quality datasets necessary for world-class model performance; • secure and maintain relationships with cloud, ecosystem, and other partners to facilitate the diffusion of Claude;

• build, operate, and maintain the infrastructure required to deploy frontier models cost-effectively and at scale; • translate model capabilities into products that meet evolving customer needs and achieve meaningful adoption; • allocate substantial financial resources efficiently across research, infrastructure, and operations; and • navigate an evolving, uncertain, and at times divergent global regulatory landscape and engage with legislators, regulators, and other stakeholders to promote the development of regulatory frameworks that we believe are good for society, our customers, and the long-term health of the global economy. Our failure to achieve any of the foregoing could prevent us from maintaining our competitive position, cause us to fall behind competitors, reduce customer adoption of, or demand for, our models, platform, and applications, and impair our revenue growth. Limitations to our revenue growth could, in turn, constrain our ability to fund the substantial and increasing expenditures required for frontier intelligence development, including investments in compute, model training, data acquisition, and safety and security systems, as well as our ability to attract, retain, and appropriately compensate highly skilled technical and operational personnel in a highly competitive market. Any such constraints could slow our rate of development and innovation, degrade model performance or reliability relative to competitors, limit our participation in industry standards-setting and regulatory processes, and reduce our attractiveness to enterprises and prosumers. These effects could create a self-reinforcing or compounding cycle in which diminished growth limits our capacity to invest in the resources necessary to remain competitive, which in turn could further weaken our ability to sustain or increase our revenue growth, making it increasingly difficult and costly to regain technological leadership and operate at the frontier of AI. Any of the foregoing could adversely affect our business, financial condition, and results of operations. Our ability to maintain our position at the frontier of AI depends in part on our use of AI systems to accelerate our research processes, a recursive dynamic we refer to as “closing the loop,” in which successive versions of our models assist in building the next, with the goal of compounding capability gains over time. Our models have been and are increasingly becoming critical tools in designing, training, testing, and optimizing next-generation AI systems. If we fail to direct sufficient compute and research efforts to AI-assisted development, or if our competitors more effectively leverage AI to enhance their own innovation cycles, we risk falling progressively further behind as competitors continuously use AI to further

enhance their own development capabilities, enabling the creation of increasingly advanced systems at an accelerating pace. As our reliance on our own AI systems to write production code, conduct research, and build subsequent model versions deepens, any systemic flaw, bias, or error in our current models could propagate into the next versions of models and products built using those systems, compounding across successive versions of Claude models in ways that may be difficult to detect or remediate. If we are unable to manage the foregoing risks, our ability to deploy the most advanced, capable, and useful AI models could be adversely affected. In addition, our increasing use of AI systems to accelerate our own research and development may cause the capabilities of our models to advance more rapidly than our ability to develop, validate, and implement corresponding safety, alignment, and security measures. Much of our evidence regarding the pace of such AI-driven acceleration is derived from lagging indicators, and our ability to forecast or measure the rate at which our capabilities are advancing, including any acceleration attributable to our most recent or unreleased models, is limited. As a result, we have in the past not identified, and may in the future not identify, a significant acceleration in our rate of development until after it has occurred. For example, our use of AI systems has already been accelerating our research and development, and our internal estimates indicate that we may experience a further increase in the rate of progress in AI capabilities. Because the AI research and development capability threshold under our Responsible Scaling Policy applies where we “observe or expect” a specified degree of accelerated progress, we made a determination that we have met this threshold based on what we expect from future models and our overall rate of progress, rather than on what we have observed from current models. We continue to assess the underlying data. We are also evaluating the implications of meeting the AI research and development capability threshold and will release further reports in the future. If our capabilities were to outpace our safety, alignment, and security measures, or if we were unable to anticipate the pace of such advancement, we may be unable to adequately assess or mitigate the resulting risks, which could increase the likelihood of safety incidents, regulatory intervention, or reputational harm, and could adversely affect our

business, financial condition, and results of operations. In light of these uncertainties, including limits on our ability to observe the pace at which our capabilities are advancing, we have adopted, and may continue to adopt, precautionary safety, alignment, and security measures to allow our safety efforts to advance together with model capabilities. In the past these measures have included delaying the release of models and limiting their initial availability, such as with Mythos through Project Glasswing, and in the future, as we learn more about model capabilities and behavior, they could include additional measures that affect whether, how, and when we develop or deploy new models or capabilities. We cannot predict with precision when we will determine that such measures are necessary, what measures will be required, or how long they will remain in place, and as the capabilities of our models advance, such measures may become more frequent and more costly. These measures could slow our capability gains relative to competitors that do not adopt similar constraints, and could adversely affect our competitive position, revenue growth, and results of operations. The development of frontier models requires substantial investments of compute resources, capital, and human capital to conduct large-scale training runs that may span weeks or months and cost many millions of dollars. There is no guarantee that any given training run will succeed—a run may fail for a number of reasons, such as failing to achieve the targeted capability or safety thresholds, producing unexpected or undesirable behaviors that cannot be adequately remediated, or yielding a model that is ultimately not competitive with alternatives available in the market. We would also face the additional cost and delay of initiating a new training run to attempt to achieve the intended model improvements. We have in the past experienced, and may in the future experience, such failed or underperforming training runs, resulting in a loss of invested resources with limited or no corresponding benefit to our business. Given the scale of compute and other resources required for frontier model training and the extended timelines involved, a failed or underperforming training run could also delay our product roadmap, affect our competitive position, and adversely affect our business, financial condition, and results of operations. New models, features, or capabilities may not be introduced on schedule, may not perform as intended, may require significant iteration after release, may fail to gain meaningful adoption, or may take longer than expected to gain meaningful adoption, any of which could reduce customer demand or place us at a competitive disadvantage. In addition, considerations relating to model safety, interpretability, alignment, and welfare may lead us to limit or delay the deprecation of older models or to continue supporting legacy versions alongside newer releases, creating tradeoffs in product development and compute distribution that could adversely affect innovation velocity, infrastructure utilization, and revenue. Customers who have integrated older models into their workflows may resist migration to newer versions, and any accelerated or forced deprecation could result in customer attrition, contractual disputes, or reputational harm. We

have also changed, and may in the future change, usage limits, rate limits, or model and product capabilities, which could adversely affect our reputation and the perceived value of our offerings. Some of these changes have resulted, and could in the future result, in claims that our disclosures regarding model access or capabilities were inadequate or misleading. Further, the industry benchmarks commonly used to test AI models may not accurately reflect real-world performance or utility, and even models that perform well on such benchmarks may not meet customer expectations when deployed in production environments, which could lead to customer dissatisfaction, reduced usage, or reputational harm. If we are unable to effectively innovate and improve our models, platform, and applications while addressing these challenges, or if competitors are able to deliver AI solutions that are perceived as more capable, safer, or more reliable, or that offer adequate performance at a lower cost, our business, financial condition, and results of operations could be adversely affected. In addition, innovation in AI involves addressing complex and evolving challenges related to data quality, data integrity, training methodologies, bias, safety, accuracy, and the behavior of model outputs. AI algorithms may be flawed, and AI systems may contain bugs, errors, or defects or fail to follow instructions or produce inaccurate, misleading, offensive, inappropriate, or otherwise harmful outputs; reflect or amplify unwanted or unknown biases present in training data; or raise concerns relating to intellectual property, privacy, or other legal and ethical considerations. We must continually invest resources to identify, mitigate, remediate, and reduce the occurrence of such issues, including through managing and improving datasets, training techniques, evaluation methods, safeguards, and human oversight. However, we may not be able to obtain the datasets that we need, including through licenses or strategic data partnerships. In addition, we or our data providers may use synthetic or AI-generated data in the training process, which could introduce quality issues, amplify biases or errors, contribute to model degradation, or raise novel legal questions regarding ownership, licensing, or

intellectual property rights that may not be timely or clearly resolved by courts or regulators. Such unintended consequences, failures, or misuse of our models could impair customer trust, slow adoption, expose us to legal or regulatory scrutiny, or adversely affect our brand and reputation. If we are unable to secure adequate compute or if we fail to direct our compute effectively across our workloads, we may be unable to train models that remain competitive with those of our competitors or adequately serve our customers. Our ability to develop, train, improve, and deploy frontier models depends on our access to enormous and increasing quantities of specialized compute. The training of Claude models has required, and we expect will continue to require, substantial compute. Industry participants, academic researchers, and our own experience suggest that, in certain contexts, applying greater amounts of compute during model training has been associated with improvements in model capabilities. These observations are consistent with aspects of broader empirical relationships among model performance, model size, training data, and compute scale that are commonly referred to as “Scaling Laws,” although the extent, efficiency, and durability of such improvements are uncertain, and these relationships may plateau, diminish, or fail to materialize at larger scales or under different conditions. Accordingly, our ability to remain at the frontier of AI capabilities is substantially dependent on our ability to secure compute at the scale required to sustain all our workloads, including to train next-generation models, and we cannot provide assurance that increased investment in compute will result in commensurate improvements in model performance. The market for compute is intensely competitive. We compete for access to compute with some of the world’s largest and most well-capitalized technology companies, many of which develop their own competitive AI models and have significantly greater financial resources than us. Many own and operate cloud computing data center infrastructure, maintain captive chip manufacturing or design capabilities, have long-standing strategic relationships with key hardware suppliers, or have significant equity investments from, or other financial arrangements with, such suppliers. These competitors can outbid us for scarce compute, secure priority allocations from chip manufacturers and CSPs, and make long-term capital commitments that we cannot match. We expect demand for advanced compute to continue to grow, and the cost of compute could rise significantly, particularly if AI-driven demand accelerates or if the ability to build highly capable models becomes more broadly available and shifts competition toward access to advanced compute. Additionally, numerous well-funded AI companies are pursuing large-scale AI model development, further increasing demand for limited compute. Certain CSPs from which we obtain substantial compute also develop and offer their own AI models and related services. As a result, our commercial relationships with these providers may not always be fully aligned with our long-term interests, and if the compute we have access to from third parties is curtailed, repriced, or terminated, whether for competitive reasons or otherwise, our business, financial condition, and results of operations could be adversely affected. In addition to competition from third parties, our ability to secure compute is subject to

allocation decisions by our CSPs. Although our CSPs are generally subject to contractual commitments with us regarding the provision of compute, our ability to obtain such capacity remains dependent on those providers' operational execution, adherence to the terms of such contractual provisions, and broader business considerations, such as capacity constraints, internal demand for compute to support their own AI and computing operations, relationships with or contractual obligations to third parties (which may include other AI companies), and risk management considerations. As a result, we may not receive the full amount of compute we seek, or such compute may not be delivered on the timelines we require, even where contractual arrangements are in place. Additionally, we have in the past sought quantities of compute from CSPs in excess of what those partners were willing to provide. If our CSPs reduce, limit, or fail to increase our allocations to keep pace with our growing requirements for any reason, we may need to seek alternative compute sources on less favorable terms or timelines, and we may be unable to secure adequate alternative compute, or any at all. In addition, certain of our compute providers and infrastructure partners depend on third-party financing or credit support to acquire chips and develop capacity, and if such financing or support becomes unavailable or more costly, contracted capacity may be delayed, repriced, reduced, or not delivered. Any such shortfall could constrain our ability to train models, serve customers, or maintain our competitive position. Even if we are able to secure compute, we must still direct it efficiently across the full range of workloads that are a part of our business, including research, model development, serving customers, and internal usage of Claude. While we manage our compute as a single, fungible pool and direct available compute between

workloads based on where the marginal unit is expected to be most productive at any given time, there can be no assurance that this approach to compute will be the most efficient or that it will support the needs of each of our workloads. Because demand for compute has historically exceeded available capacity, any constraint on our ability to procure, maintain, or expand that capacity could adversely affect multiple aspects of our business simultaneously, including our ability to serve customers at the quality and scale they expect, to train and improve our models, and to develop new products and capabilities. While we actively manage and balance such compute capacity across these workloads via our capacity management system, if this system fails to perform as intended, or if our ability to dynamically rebalance capacity across workloads is disrupted, customer-facing service quality or model and product development timelines could be adversely affected. There can be no assurance that we will make optimal resourcing decisions, and any significant inefficiencies, including those related to deployment and rebalancing, could adversely affect both our near-term results of operations and our long-term strategic position. These compute resourcing challenges are heightened by the growth of our agentic offerings, which involve multi-step reasoning chains, iterative tool invocations, extended context windows, and sustained inference sessions that consume significantly more compute per user interaction than single-turn API calls or chat-based interactions. As agentic workloads grow as a proportion of our total compute demand, the aggregate compute required to serve our customers could increase disproportionately relative to the associated revenue, placing additional strain on our compute workload management decisions. The majority of our compute is procured under multi-year commitments that are typically entered into well in advance of the periods in which that compute is deployed. Because different accelerators have varying performance characteristics across workload types, the composition of our committed compute portfolio, including the mix of accelerator architectures procured from our partners, reflects our demand forecasts and development roadmap assumptions at the time those commitments are made. If the magnitude or mix of workloads we need to support differs materially from those forecasts, whether due to changes in model architecture, customer demand, internal operations, or research priorities, the composition of our available compute may not optimally match our needs in a given period. This could affect our ability to serve customers, develop models, and support internal operations. If we fail to secure sufficient compute, or if we fail to effectively direct the limited compute we have across our workloads to develop and deploy models with capabilities at least comparable to those of our competitors, we risk falling behind the technological frontier. The pace of AI development is extremely rapid, with successive versions of models emerging in compressed timeframes, often measured in months rather than years. As a result, even a brief period without access to sufficient compute could cause us to fall behind quickly. With AI development, the gap between frontier models and those one or two versions behind can be significant in terms of capability, performance, and commercial applicability. Customers, partners, and software developers may shift to competitors whose models offer superior capabilities, and we may be unable to attract or retain the technical talent necessary for frontier research if our researchers lack access to adequate compute or if we are not perceived as operating at the forefront of model development. Any inability to maintain our position at the frontier could result in loss of customers and revenue,

diminished brand and reputation, and a decline in our competitive position, each of which may be difficult to regain. If we are unable to manage the foregoing risks, our business, financial condition, and results of operations could be adversely affected. If our actual compute requirements differ materially from our projections, we may fail to secure sufficient compute to meet our needs or we may be committed to excess compute that remains underutilized. The competitive and structural dynamics of the compute market require us to make binding financial commitments for compute months or years in advance of our day-to-day delivery and use. Chip manufacturers, CSPs, and data center developers require advance commitments in order to allocate production capacity, plan infrastructure investments, and manage their own supply chains. As a result, we must forecast our compute requirements well in advance of their deployment and commit capital accordingly, often with limited ability to adjust commitments if our needs change. These forecasts for future compute usage are inherently uncertain. Our actual compute requirements depend on numerous factors that are difficult to predict, including the pace of our research progress, the compute intensity of training techniques we develop or adopt, changes in the competitive landscape that may require us to accelerate or modify our development plans, customer demand for inference compute, and the evolution of hardware or other technological efficiencies and changes. If we underestimate our compute requirements, we

may be unable to secure sufficient compute to support our forecasted customer demand and model development roadmap. In the past, we have significantly underestimated our compute requirements, which left us capacity constrained and impacted both product development and deployment. Further, the timing of compute coming online may not align with the timing of demand growth, which could result in periods of excess or underutilized compute and committed cash outflows that do not produce immediate revenue. Our ability to forecast and secure the compute we need is further challenged by our ongoing transition from a cloud-only, fixed-time arrangement compute model to a hybrid model that includes directly leased chips that are hosted in dedicated data center facilities leased and operated on our behalf, to our specification. This transition introduces additional variables into our procurement planning, including the timing of facility completion, the pace at which newly commissioned facilities achieve full operational capacity, our ability to operate, maintain, and integrate on-premise capacity with cloud-based infrastructure, the operational execution of our suppliers, and the performance of hardware in new facility environments. Delays in securing energy, water, and other critical services for these facilities could further slow the pace at which newly commissioned compute becomes available for use. Errors in forecasting any of these transition-related variables could compound our risks related to securing sufficient compute and result in either capacity shortfalls or underutilized capacity during the transition period. We have entered into, and expect to continue to enter into, long-term agreements for compute and related resources that involve minimum purchase commitments, take-or-pay provisions, and other terms that obligate us to pay for resources regardless of actual utilization. These commitments represent substantial fixed obligations. As of and subsequent to September 30, 2026, we have entered into purchase commitments and lease agreements, including short-term commitments, as well as leases not yet commenced, with an aggregate committed spend of approximately $900.9 billion , substantially all of which is related to compute and infrastructure. These commitments include approximately $451.6 billion expected to be incurred before December 31, 2030, and substantially all of the remainder is expected to be incurred between 2031 and 2045. Approximately $70.7 billion of these commitments are due in the 12 months following September 30, 2026. See the section titled “Management’s Discussion and Analysis of Financial Condition and Results of Operations— Liquidity and Capital Resources—Contractual Obligations” for more information. Our ability to meet the substantial and growing financial obligations of these requirements to secure compute depends on our ability to generate significant revenue growth. We must grow our customer base and expand revenue rapidly in order to generate the cash flows necessary to satisfy these commitments. However, our revenue is primarily consumption-based, and much of it is not subject to long-term contractual commitments, while our compute and infrastructure commitments generally extend over multiple years. As a result, we may enter into these obligations based on revenue expectations that may not be realized. If our revenue growth does not keep pace with our compute-related payment obligations, if we are unable to successfully convert our compute investment into revenue-generating customer workloads, if we fail to raise additional capital on acceptable terms, or if we experience other adverse developments affecting our liquidity, we may be unable to satisfy these obligations. Any such failure could also result in defaults under our agreements, loss of access to critical resources, and reputational harm. In particular, a default under certain commercial agreements could trigger cross-default provisions under the A&R Credit Agreement (as defined below) and certain of our outstanding convertible promissory notes, potentially resulting in the acceleration of substantially all of our outstanding, and

any future, indebtedness. Such cross-defaults could set off a cascading series of adverse consequences, including the exercise of remedies by multiple creditors, the termination of key commercial arrangements, enforcement actions against our assets, and severe liquidity constraints. An acceleration of our indebtedness, combined with the loss of access to critical data center infrastructure, among other things, could impair our ability to train and deploy models, serve customers, and fund ongoing operations, which could have an adverse effect on our business, financial condition, and results of operations. For more information, see the sections titled "—We have incurred significant indebtedness in order to operate and grow our business, including to fund our substantial compute commitments, and our ability to service this indebtedness may be constrained by our cash flows and liquidity" and "—We may not have the ability to raise the funds necessary to repay the principal and accrued interest of our outstanding convertible promissory notes or the A&R Credit Agreement in cash when required by the terms of the notes or the A&R Credit Agreement, as applicable, or at their maturity."

Additionally, the rapid pace of technological change across AI infrastructure, hardware, software, system architectures, and model development techniques means that compute secured today may become less efficient or economically attractive relative to future alternatives. We may find ourselves committed to using older-generation hardware and infrastructure while competitors gain access to more advanced systems, or we may be paying premium prices for resources that have declined in market value. These dynamics could adversely affect our cost structure and competitive position. Certain of our financing arrangements require us to enter into exclusive or preferential supply commitments with particular chip suppliers with respect to the compute financed pursuant to such arrangements. These commitments may limit our ability to shift procurement to alternative suppliers or newer chip architectures if our needs or the technology landscape evolve, potentially leaving us committed to platforms that are suboptimal for our actual workload requirements and reducing our ability to adjust our compute portfolio in response to changed circumstances. If we are unable to manage the foregoing risks, our business, financial condition, and results of operations could be adversely affected. Constraints on the critical inputs supporting compute could limit our ability to secure adequate compute. The compute necessary for training and deploying frontier models depends on a supply chain with multiple potential points of constraint or failure. The availability, pricing, and lead times for compute are affected by the supply of advanced chips and other IT equipment, the availability of sufficient electrical power and the infrastructure required to generate and deliver it, and the capacity of data center infrastructure. Each of these inputs faces significant constraints that may limit our ability to secure the compute we need. In addition, increasing demand for compute has contributed to, and may continue to contribute to, increases in the cost of such resources, which could limit our ability to invest in, develop, or deploy our models, platform, and applications at the pace we anticipate. Supply and Concentration of Chips, Memory, and Data Center Equipment The specialized chips required for AI workloads, particularly high-performance accelerators and memory, are produced by a limited number of manufacturers. The production of advanced chips is extraordinarily complex, requiring manufacturing processes at the leading edge of technological capability. The majority of the world's most advanced chip manufacturing capacity is concentrated in Taiwan, primarily at Taiwan Semiconductor Manufacturing Company ("TSMC"). This geographic concentration creates significant risks, including exposure to natural disasters such as earthquakes and typhoons, potential disruption from geopolitical tensions between Taiwan and China, vulnerability to any interruption in Taiwan's power grid, water supply, or other critical infrastructure, and susceptibility to developing trade regulations and restrictions. Any military conflict, tariffs and trade restrictions, economic coercion, blockade, natural disaster, or other incident affecting Taiwan could severely disrupt the global supply of advanced chips, with effects that could persist for years given the time required to construct and qualify alternative fabrication facilities. Efforts to establish sufficient advanced chip manufacturing capacity outside of Taiwan, including in the United States, remain in early stages and may not produce adequate compute to materially reduce these concentration risks within the foreseeable future. In addition, due to the demand for advanced chip manufacturing, TSMC has, at times, been unable to fully satisfy market demand for chip manufacturing capacity, and as a result, we have not always been able to obtain sufficient quantities of chips on a timely basis. Accordingly, any restrictions or delays on the supply of advanced chips out of Taiwan or TSMC would adversely affect our business, financial condition, and results of operations. In addition, both high-bandwidth memory, which is a key component in high-performance accelerators, and other high-performance memory products used in compute infrastructure, are largely manufactured by SK hynix Inc., Samsung Electronics Co., and Micron Technology, Inc. Each of these manufacturers is working to expand its production capacity to meet demand. However, there is no guarantee that these efforts will result in adequate supply fast enough to meet demand, particularly if large customers or competitors secure exclusive supply

commitments or otherwise reserve substantially all of a manufacturer's available production capacity. The concentration and constraints in manufacturing supply expose us to the risk that manufacturers may make allocation decisions or strategic shifts that adversely affect our access to the advanced chips needed for our business. We also face risks that specialized chips may contain defects, or that their production or delivery may be delayed or disrupted, which, given long replacement lead times, could constrain available compute and

adversely affect our operations and ability to scale. A meaningful portion of the components we and our partners rely on is sourced from or manufactured outside of the United States, and new or increased tariffs, import restrictions, or trade disputes could delay deliveries, force supply-chain changes, or raise our costs, some of which we bear directly under pass-through arrangements. For more information, see the section titled "— Governmental trade controls, including export and import controls, sanctions, customs requirements, and related regimes, could subject us to liability or loss of contracting privileges or limit our ability to compete in certain markets." We also depend on third-party chip design and services companies for the design, manufacture, and support of certain custom accelerator platforms, as well as for related infrastructure financing arrangements. Disruption to these relationships, or any failure by these parties to deliver hardware or services on expected timelines and specifications, could adversely affect the availability, performance, or cost of accelerator platforms on which we rely. Data centers also require electrical, generation, and transmission equipment, cooling systems, backup generators, and networking hardware, some of which are made by only a few manufacturers, must be ordered years in advance, and require substantial pre-delivery payments that we may not recover if our plans change. In addition, certain equipment, such as high-voltage transformers, switchgear, turbines, generators, substations, and transmission components, has been subject to extended delivery timelines, price increases, tariffs and national security requirements, and allocation constraints as demand from data center operators and other large electricity consumers has grown. If such equipment is delayed or unavailable, we may be able to replace it only at higher prices or under less favorable long-term commitments. In order to secure near-term access to compute, we may also elect, or be required, to procure long-lead-time generation and electrical equipment at prices above prevailing market rates on the expectation that it can be deployed on an accelerated timeline to unlock capacity in the near term. Realizing the benefit of paying such prices depends on the timely and successful engineering, procurement, and construction of the associated facilities, which requires specialized expertise that is currently in short supply and involves significant execution risk. Delays, cost overruns, or failure to secure energy on the anticipated schedule could mean that we do not obtain the energy and capacity for which we paid such prices. In some cases, these arrangements also depend on a developer or other counterparty agreeing to purchase the equipment at such prices, finance such purchase, and recover those costs and associated financing costs from us over time, and we may bear these costs even if the underlying capacity is delayed or never realized. Our equipment supply framework agreements and similar arrangements may expose us to further risks, including minimum purchase or take-or-pay commitments, price escalation, and obligations that we are unable to align with our actual deployment needs. Additionally, the U.S. government has imposed and may continue to impose export controls restricting the sale of advanced chips and chip manufacturing equipment to certain countries, including China. These restrictions have affected global chip and related equipment supply allocation, and future changes in export control policy could affect the availability or pricing of the compute resources on which we rely. Future export controls could be more restrictive, for example, by placing restrictions on the countries in which advanced chips may be installed or in which model weights may be stored, which could limit where we or our partners are able to build data centers, constrain our international expansion, or increase our costs. Electrical Power and Fuel Constraints The operation of data centers at the scale required for frontier model development and deployment consumes substantial quantities of electrical power. The availability of sufficient electrical power has emerged as a significant constraint on data center development and expansion in many geographies, including the United States. Electric grid infrastructure in the United States and other countries was not designed to accommodate the rapid increase in demand associated with AI development, and the lead times required to build new generation capacity, transmission infrastructure, and grid interconnections can extend for years. Demand on electric grid infrastructure is also increasing as a result of the increased electrification of homes, vehicles, and industry, and this growing demand is occurring at a time when power generation capacity additions, transmission expansion, and grid modernization efforts have not kept pace and certain existing generation resources are being retired, all of which has contributed to higher electricity prices and longer interconnection timelines in certain regions. We have publicly committed to covering the cost of electricity price increases to consumers caused by our data centers, including both the cost of electricity itself and the

infrastructure necessary to deliver it.

Our access to electrical power also depends on our ability, and the ability of the third parties on which we rely for the construction, management, and operation of our data center facilities, to negotiate and maintain contracts with utility providers and grid operators and to obtain approvals from public utility commissions and other regulators. These parties may be unable or unwilling to approve, or may condition or delay, the connections, capacity expansions, special or negotiated rate arrangements, and infrastructure investments necessary to support data center growth, or may grant them only at costs that limit the economic efficiency and viability of certain facilities, such as by requiring us to bear grid upgrade or interconnection costs beyond those attributable to serving our facilities—representing a larger share of costs than is required under existing or prior regulatory constructs. Interconnection queues for large new electricity loads have also lengthened substantially in many regions, which could delay or limit our access to grid power. Utility providers and regulators may also reallocate previously committed capacity that our data center facilities may require. In many jurisdictions, a single utility provider has substantial control over the terms, timing, and cost of the service and infrastructure necessary to power our data center facilities. The commercial and regulatory priorities of utility providers may differ from ours, including with respect to which infrastructure investments are made, how their costs are allocated, and whether alternative or self-supplied generation arrangements are permitted and on what conditions. Utility providers and other stakeholders may also seek regulatory or legislative measures affecting large electricity loads, self-supply, or alternative power generation arrangements. The outcomes of our negotiations and regulatory proceedings involving utilities are uncertain, and unfavorable results could increase our costs, delay our projects, or limit the alternatives available to us. More broadly, regulatory proceedings relating to cost allocation, rate design, and the treatment of large electricity loads are ongoing in numerous jurisdictions, and adverse outcomes, changes in policy, or the imposition of new tariffs, surcharges, or curtailment obligations applicable to data centers could increase our costs or limit our access to electrical power. In addition, the construction of the generation and transmission infrastructure on which we depend is itself capital-intensive and exposed to labor shortages, supply chain disruptions, permitting delays, changes in interest rates, tariffs, other cost overruns, and financing constraints. The significant and increasing capital required to fund these generation, transmission, and related infrastructure investments, whether incurred by us, our utility and energy supply providers, or our data center development partners, is sensitive to prevailing interest rates and the availability of debt or equity financing on acceptable terms. Any inability of our partners to access capital when needed, and the limitations, cost, and competition for such capital, could increase the cost of the power on which we depend, delay or halt projects, or render certain facilities commercially unviable, and under certain of our arrangements, we bear these financing and capital costs directly or on a pass-through basis. Furthermore, newer sources of generation, such as solar, wind, battery storage, nuclear, and geothermal power, may not be available in the quantities or on the timelines we require, or may be available only at higher cost. Our facilities separately depend on access to reliable water for cooling, which in some regions is scarce, subject to use restrictions, or increasingly costly. Any of the foregoing could adversely affect the availability of the power or water we require, increase our costs, or limit our ability to develop, expand, or operate our data center facilities, which could adversely affect our business, financial condition, and results of operations. Competition for access to electrical supply and data center locations with adequate power availability is intense, and we compete for such access with well-capitalized CSPs, some of which produce and develop competing AI models, cryptocurrency mining operations, and other large-scale electricity consumers. In addition to, or sometimes in lieu of, grid-based electricity, an increasing share of data centers are being developed with dedicated on-site power generation capacity that supplies electricity directly to the facility independent of any utility-operated electric grid. The reliability of such off-grid power generation solutions varies by the underlying technology, the equipment used, and the operator and will reflect site-by-site operational and environmental profiles. Development of this new capacity in turn requires power generation and electrical infrastructure that are themselves subject to supply challenges and operational complexity. For example, certain of our planned data center facilities rely on dedicated on-site natural gas-fired generation, which exposes us to fuel price volatility and supply disruptions that are outside of our control and for which we bear the cost on a pass- through basis. These facilities are also subject to constraints on pipeline capacity and firm fuel transportation necessary to deliver natural gas. As data center operators increasingly rely on natural gas-fired generation, and because behind-the-meter and on-site arrangements are feasible only in some jurisdictions given the applicable regulatory frameworks, we expect increased competition among users for the same natural gas supply and pipeline capacity in those areas. This competition could drive up the price of natural gas in such areas

and may

adversely affect our ability to contract for firm fuel supply and transportation and to secure the pipeline connections and laterals our facilities require. These dynamics present several related risks. We may be unable to secure firm gas supply and transportation during the development of a facility, causing a planned site to be delayed or abandoned; we may face fuel supply or delivery disruptions once a facility is operating; and because our arrangements generally provide for fuel and related costs to be borne by us on a pass-through basis without cost caps, we remain exposed to increases in the cost of natural gas and its transportation. More broadly, the cost of electricity is subject to significant variability that is largely outside of our control, as wholesale electricity prices fluctuate with changing supply and demand, weather, fuel prices, tariffs or surcharges, and grid stress, and can rise sharply during periods of peak demand or constrained supply, and such increases in energy costs could adversely affect our results of operations. In some regions, including certain regions where we have contracted compute commitments, data center development has prompted public controversy and regulatory scrutiny regarding environmental impact, land and water use, and the allocation of limited power resources. For example, if our efforts to secure electrical power are perceived to affect other utility customers and the broader grid, regulators may impose conditions, standby charges, or other requirements intended to address such impacts. More broadly, our access to electrical power is subject to an evolving regulatory framework characterized by complexity, jurisdictional overlap, and rules that were largely designed around service by incumbent utility providers and existing generation arrangements and that may not readily accommodate new entrants or alternative approaches, which can result in significant delay or prevent the development of the generation we require. Grid-tied electrical supply, for instance, is subject to the jurisdiction of federal and state energy regulators, regional grid operators, and public utility commissions, while behind-the-meter and on-site generation is subject to a distinct set of exclusive service franchise protections, permitting, environmental, air quality, interconnection, and siting requirements. In addition, environmental standards applicable to power generation continue to evolve, including as a result of changing governmental priorities and shifting public and political views regarding climate change and the appropriate treatment of various power generation sources. Any failure to secure access to adequate electrical power, any legal, regulatory, or other administrative action or change limiting the availability of electricity, or any significant increase in energy or compliance costs, could adversely affect our cost structure and limit our ability to scale our compute infrastructure, develop and deploy our models, and service our customers. We have also made, and may in the future make, voluntary commitments to offset the impact of our data center energy consumption on local communities and electricity ratepayers, including commitments to cover grid infrastructure upgrade costs, procure new power generation to match our energy needs, invest in curtailment systems to reduce strain on the grid during peak demand periods, and address environmental impacts. These commitments have increased, and we expect will continue to increase, our operating costs. Moreover, we believe this is part of a broader emerging trend in the AI industry, and such commitments, or similar obligations, may become increasingly standard or required, which could further increase our costs. In addition, the physical effects of climate change may increase our operating costs over time, as higher average temperatures can increase the power consumption and water usage required to cool equipment at the facilities on which we depend. Data Center Infrastructure Even where chips and electrical power are available, the development of data center facilities requires substantial time, capital, and specialized expertise. Data center construction involves lengthy permitting processes, complex supply chains for specialized equipment, and significant lead times for the installation and commissioning of cooling systems, electrical infrastructure, and networking equipment. Labor shortages, including shortages of, and competition for, electricians, facility technicians, and other skilled tradespeople, supply chain disruptions, permitting delays, and other factors can extend construction timelines and increase costs. Further, given data center development projects require significant advance capital commitments with dedicated end users, we have in some cases relied, and expect to continue to rely, on arrangements for financing or with partners who guarantee the lease obligations to support the development of data center capacity that is expected to host such partner’s chips or computing infrastructure. In addition, there are a limited number of qualified third-party vendors with the expertise we require and that support the particular infrastructure we wish to implement. Accordingly, we rely on this limited number of third parties for the construction, management, and operation of certain of our data center facilities, some of which have limited operating history or may have less experience operating data center infrastructure at the scale and complexity our operations

require. In certain cases, these third-party operators are trained or supervised by our chip partners to manage facilities on our behalf but are not affiliates of those partners and may not have the same depth of expertise, operational resilience, or resources. Any failure by these third-party operators to perform as expected could disrupt our access to compute. In addition, safety or environmental incidents at facilities built or operated for us, whatever their cause, could expose us to claims, fines, or suspensions of construction or operations, the costs of which we may bear. We also depend on an extended network of suppliers, builders, and facility operators. If any of them engage in practices that are unsafe, unlawful, or contrary to the labor, human rights, environmental, or other standards we or our partners require, we could face legal liability, lose the ability to use affected suppliers or facilities, and suffer harm to our reputation, even where we did not direct or know of the conduct. Our collective data center infrastructure is currently concentrated primarily in the United States, which exposes us to risks associated with regional regulatory changes, power grid constraints or failures, disruptions to network connectivity between facilities, natural disasters, physical attacks—including those motivated by geopolitical events, opposition to AI, and opposition to data center development—and other disruptions that could affect multiple facilities. Certain state, local, and municipal jurisdictions in the United States are also considering or have approved moratoriums on new data center construction, citing public health, safety, and environmental concerns. Such jurisdictions and other regulators, grid operators, or utilities may also adopt other measures that subject users of large electricity loads, including data centers, to curtailment, interruptions, forced reductions, disconnection, or other restrictions during periods of grid stress, whether through changes in law or policy, market or reliability constraints, limits on transmission and distribution infrastructure, or a preference for protecting other utility customers and the broader grid. New large loads on the electric grid may be increasingly subject to mandatory curtailment in regions where there is insufficient grid-connected power generation to reliably supply existing and new electricity loads. For example, states such as New York and Texas have recently moved to regulate large-scale data centers, whether by pausing state permitting for new large facilities, as New York has, or by subjecting large electricity loads to curtailment or disconnection during grid emergencies, as Texas has. Such moratoriums, or similar actions, could delay or prevent the construction or expansion of facilities on which we depend for compute. Opposition to data center development has also become more organized and coordinated across jurisdictions, and our public profile may draw particular attention to projects associated with us. The loss of, or a significant disruption at, any major facility may not be readily absorbed by our remaining or future infrastructure. We are expanding our data center capacity internationally as part of our efforts to diversify our infrastructure footprint and reduce concentration risks associated with operating a significant portion of our compute infrastructure in the United States and to support our international growth strategy, including by entering into letters of intent and negotiating lease and related agreements for data center facilities in multiple non-U.S. jurisdictions, which exposes us to additional risks related to foreign regulatory requirements and approvals, data sovereignty and localization laws, governmental access or assistance requirements, cross-border data transfer restrictions, foreign tax regimes, and increased operational complexity. Our ability to enter into definitive agreements for data center capacity and operations in non-U.S. jurisdictions depends on the timely formation, registration, and receipt of governmental approval of local subsidiaries, as well as the establishment of local banking arrangements, and any delay or failure in completing these steps could delay or prevent us from mitigating our infrastructure concentration risks, executing our international growth strategy, securing planned compute capacity, or could result in substantial additional taxes, penalties, or other liabilities. In addition, regulatory, legal, and other constraints in many jurisdictions may limit the locations in which we can practically site or contract for data center capacity, which could constrain our ability to diversify our infrastructure footprint as quickly or as broadly as we currently intend. We and our principal competitors rely on substantially similar inputs to develop and deploy frontier models, including advanced chips sourced from a small number of suppliers and data center capacity and compute obtained from CSPs and other third parties. As we and our competitors increasingly pursue vertical integration, the basis of competition could shift over time to encompass not only which company can develop the most capable models, but also which company can develop the best underlying components. Our future competitiveness may depend in part on our ability to successfully develop vertically integrated compute infrastructure on competitive terms and timelines, and there can be no assurance that we will be able to do so. Our competitors may be better positioned, or may invest greater efforts, to vertically integrate across the compute supply chain, including through proprietary chip design, fabrication, ownership or long-term control of data center infrastructure, access to dedicated power generation assets, exclusive supply arrangements with

wheresthe.ai source copy, page 46 of 47

hardware manufacturers, or by securing the limited sites suitable for large-scale data center development ahead of us. Competitors that succeed in controlling a greater share of these supply chain layers may be able to secure compute at greater scale, lower cost, or with greater reliability than we can, particularly during periods of high demand or supply constraint, when captive supply relationships may confer certain advantages over participants that depend on third-party providers. Such competitors may also develop proprietary hardware- software co-designs that optimize compute efficiency for their specific model architectures, which could compound any capability gap across successive model versions. While we are pursuing initiatives to develop dedicated data center infrastructure, there can be no assurance that these efforts will be completed on the timelines or at the costs we anticipate, or that they will be sufficient to offset any structural advantages that vertically integrated competitors may achieve. Competitors that succeed in these efforts may be better positioned to remain at the frontier of AI regardless of broader market constraints. Natural hazards, extreme weather, or other disruptions affecting the regions where these inputs are produced or where our facilities are located could further constrain the availability of chips, power, water, or data center capacity on which we depend. If we are unable to manage the foregoing risks, our business, financial condition, and results of operations could be adversely affected. We are highly dependent on the continued service of our co-founders, senior management team, and key technical employees, and the loss of any of these individuals may adversely affect our business, financial condition, and results of operations. We are highly dependent on the continued service of our Chief Executive Officer, President, co-founders, senior management team, and key technical