

Amazon agrees to invest up to \$4 billion in Anthropic

Source copy kept by wheresthe.ai for the story "Amazon agrees to invest up to \$4 billion in Anthropic" (<https://wheresthe.ai/d/amazon-anthropic-4bn-investment-2023-09/>).

Source: <https://www.aboutamazon.com/news/company-news/amazon-aws-anthropic-ai>

Text saved 18 Oct 2026 (the page as published).

SHA-256 of the saved text: cc1deea6c82084d88e98bfde1f6a3c458b583113faa0f54c40dc895f83abfe56

Anthropic selects AWS as its primary cloud provider and will train and deploy its future foundation models on AWS Trainium and Inferentia chips, taking advantage of AWS's high-performance, low-cost machine learning accelerators.

Anthropic deepens commitment to AWS, making its future foundation models accessible to millions of developers and providing AWS customers early access to unique features for model customization, using their proprietary data, and fine-tuning capabilities, all through Amazon Bedrock.

Amazon and Anthropic today announced a strategic collaboration that will bring together their respective industry-leading technology and expertise in safer generative artificial intelligence (AI) to accelerate the development of Anthropic's future foundation models and make them widely accessible to AWS customers. As part of the expanded collaboration:

Anthropic will use AWS Trainium and Inferentia chips to build, train, and deploy its future foundation models, benefitting from the price, performance, scale, and security of AWS. The two companies will also collaborate in the development of future Trainium and Inferentia technology. AWS will become Anthropic's primary cloud provider for mission critical workloads, including safety research and future foundation model development. Anthropic plans to run the majority of its workloads on AWS, further providing Anthropic with the advanced technology of the world's leading cloud provider. Anthropic makes a long-term commitment to provide AWS customers around the world with access to future generations of its foundation models via Amazon Bedrock, AWS's fully managed service that provides secure access to the industry's top foundation models. In addition, Anthropic will provide AWS customers with early access to unique features for model customization and fine-tuning capabilities. Amazon will invest up to \$4 billion in Anthropic and have a minority ownership position in the company. Amazon developers and engineers will be able to build with Anthropic models via Amazon Bedrock so they can incorporate generative AI capabilities into their work, enhance existing applications, and create net-new customer experiences across Amazon's businesses.

How AWS customers are approaching generative AI with Amazon Bedrock From sports to travel to life sciences, generative AI is changing how organizations work and the experiences they're building for their customers.

"We have tremendous respect for Anthropic's team and foundation models, and believe we can help improve many customer experiences, short and long-term, through our deeper collaboration," said Andy Jassy, Amazon CEO. "Customers are quite excited about Amazon Bedrock, AWS's new managed service that enables companies to use various foundation models to build generative AI applications on top of, as well as AWS Trainium, AWS's AI training chip, and our collaboration with Anthropic should help customers get even more value from these two capabilities."

"We are excited to use AWS's Trainium chips to develop future foundation models," said Dario Amodei, co-founder and CEO of Anthropic. "Since announcing our support of Amazon Bedrock in April, Claude has seen significant organic adoption from AWS customers. By significantly expanding our partnership, we can unlock new possibilities for organizations of all sizes, as they deploy Anthropic's safe, state-of-the-art AI systems together with AWS's leading cloud technology."

An AWS customer since 2021, Anthropic has grown quickly into one of the world's leading foundation model providers and a leading advocate for the responsible deployment of generative AI. Their foundation model, Claude, excels at a wide range of tasks, from sophisticated dialogue and creative content generation to complex reasoning and detailed instruction, while maintaining a high degree of reliability and predictability. Its industry-leading 100,000 token context window can securely process extensive amounts of information across all industries, from manufacturing and aerospace to agriculture and consumer goods, as well as technical, domain-specific documents for industries such as finance, legal, and healthcare. Customers report that Claude is much

less likely to produce harmful outputs, easier to converse with, and more steerable compared to other foundation models, so developers can get their desired output with less effort. Anthropic's state-of-the-art model, Claude 2, scores above the 90 th percentile on the GRE reading and writing exams, and similarly on quantitative reasoning.

7 generative AI innovations from AWS Summit New York 2023 Attendees learn how generative AI is transforming organizations, across all kinds of industries and applications.

Today's news is the latest AWS generative AI announcement as the company continues to expand its unique offering at all three layers of the generative AI stack. At the bottom layer, AWS continues to offer compute instances from NVIDIA as well as AWS's own custom silicon chips, AWS Trainium for AI training and AWS Inferentia for AI inference. At the middle layer, AWS is focused on providing customers with the broadest selection of foundation models from multiple leading providers where customers can then customize those models, keep their own data private and secure, and seamlessly integrate with the rest of their AWS workloads—all of this is offered through AWS's new service, Amazon Bedrock. With today's announcement, customers will have early access to features for customizing Anthropic models, using their own proprietary data to create their own private models, and will be able to utilize fine-tuning capabilities via a self-service feature within Amazon Bedrock. At the top layer, AWS offers generative AI applications and services for customers like Amazon CodeWhisperer, a powerful AI-powered coding companion, which recommends code snippets directly in the code editor, accelerating developer productivity as they code.

As part of this deeper collaboration, AWS and Anthropic are committing meaningful resources that are helping customers get started with Claude and Claude 2 on Amazon Bedrock, including through the AWS Generative AI Innovation Center, where teams of AI experts will help customers of all sizes to develop new generative AI-powered applications to transform their organizations.

Customers accessing Anthropic's current models via Amazon Bedrock are building generative AI-powered applications that help automate tasks such as producing market forecasts, developing research reports, enabling new drug discovery for healthcare, and personalizing education programs. Enterprises already taking advantage of this advanced technology include Lonely Planet, a premier travel media company celebrated for its decades of travel content; Bridgewater Associates, a premier asset management firm for global institutional investors; and LexisNexis Legal & Professional, a top-tier global provider of information and analytics serving customers in more than 150 countries.

"We are developing a generative AI solution on AWS to help customers plan epic trips and create life-changing experiences with personalized travel itineraries," said Chris Whyde, senior vice president of Engineering and Data Science at Lonely Planet. "By building with Claude 2 on Amazon Bedrock, we reduced itinerary generation costs by nearly 80 percent when we quickly created a scalable, secure AI platform that organizes our book content in minutes to deliver cohesive, highly accurate travel recommendations. Now we can re-package and personalize our content in various ways on our digital platforms, based on customer preference, all while highlighting trusted local voices—just like Lonely Planet has done for 50 years."

AWS reaffirms its commitment to responsible generative AI As a pioneer in artificial intelligence and machine learning, AWS is committed to developing and deploying generative AI responsibly.

"At Bridgewater, we believe the global economic machine can be understood, so we strive to build a fundamental, cause-and-effect understanding of markets and economies powered by cutting-edge technology," said Greg Jensen, co-CIO at Bridgewater Associates. "Working with the AWS Generative AI Innovation Center, we are using Amazon Bedrock and Anthropic's Claude model to create a secure large language model-powered Investment Analyst Assistant that will be able to generate elaborate charts, compute financial indicators, and create summaries of the results, based on both minimal and complex instructions. This flexible solution will accelerate the more mundane, yet still involved, steps of our research process, enabling our analysts to spend more time on the difficult and differentiated aspects of understanding markets and economies."

"We are working with AWS and Anthropic to host our custom, fine-tuned Anthropic Claude 2 model on Amazon Bedrock to support our strategy of rapidly delivering generative AI solutions at scale and with cutting-edge encryption, data privacy, and safe AI technology embedded in everything we do," said Jeff Reihl, executive vice president and CTO at LexisNexis Legal & Professional. "Our new Lexis+ AI platform technology features conversational search, insightful summarization, and intelligent legal drafting capabilities, which enable lawyers to increase their efficiency, effectiveness, and productivity."

Amazon and Anthropic are each engaged across a number of organizations to promote the responsible development and deployment of AI technologies, including the Organization for Economic Cooperation and Development (OECD) AI working groups, the Global Partnership on AI (GPAI), the Partnership on AI, the International Organization for Standardization (ISO), the National Institute of Standards and Technology (NIST), and the Responsible AI Institute. In July, both Amazon and Anthropic joined President Biden and other industry leaders at the White House to show their support for a set of voluntary commitments to foster the safe, secure, responsible, and effective development of AI technology. These commitments are a continuation of work that both Amazon and Anthropic have been doing to support the safety, security, and responsible development and deployment of AI and will continue through this expanded collaboration.

About Amazon Web Services Since 2006, Amazon Web Services has been the world's most comprehensive and broadly adopted cloud. AWS has been continually expanding its services to support virtually any workload, and it now has more than 240 fully featured services for compute, storage, databases, networking, analytics, machine learning and artificial intelligence (AI), Internet of Things (IoT), mobile, security, hybrid, virtual and augmented reality (VR and AR), media, and application development, deployment, and management from 102 Availability Zones within 32 geographic regions, with announced plans for 12 more Availability Zones and four more AWS Regions in Canada, Malaysia, New Zealand, and Thailand. Millions of customers—including the fastest-growing startups, largest enterprises, and leading government agencies—trust AWS to power their infrastructure, become more agile, and lower costs. To learn more about AWS, visit aws.amazon.com.

About Amazon Amazon is guided by four principles: customer obsession rather than competitor focus, passion for invention, commitment to operational excellence, and long-term thinking. Amazon strives to be Earth's Most Customer-Centric Company, Earth's Best Employer, and Earth's Safest Place to Work. Customer reviews, 1-Click shopping, personalized recommendations, Prime, Fulfillment by Amazon, AWS, Kindle Direct Publishing, Kindle, Career Choice, Fire tablets, Fire TV, Amazon Echo, Alexa, Just Walk Out technology, Amazon Studios, and The Climate Pledge are some of the things pioneered by Amazon. For more information, visit amazon.com/about and follow @AmazonNews.

About Anthropic Anthropic is an AI safety and research company based in San Francisco. Our interdisciplinary team has deep experience across machine learning, physics, policy, and product. Together, we create reliable, interpretable, and steerable AI systems. Anthropic's flagship product is Claude, an AI assistant focused on being helpful, harmless, and honest. Learn more about Anthropic at anthropic.com.

Forward-looking statements This communication contains forward-looking statements that are inherently uncertain and difficult to predict, including statements regarding planned investments, anticipated business activities, and expected benefits of the expanded collaboration between Amazon and Anthropic, such as expected developments, performance, scaling, or impacts as a result of the collaboration; the reliability, safety, security, applications, or effectiveness of artificial intelligence technologies; and increased or early accessibility of new features, models, and tools. We use words such as will, believes, expects, future, should, plan, potential, continue, and similar expressions, as well as words referring to future outcomes such as train, make, accelerate, benefit, bring, collaborate, develop, become, create, enhance, expand, incorporate, build, provide, change, receive, help, use, and variations of such words, to identify forward-looking statements. Actual results and outcomes could differ materially for a variety of reasons, including, among others, changes in global economic conditions and customer demand and spending, inflation, interest rates, regional labor market and global supply chain constraints, world events, the rate of growth of cloud services, the amount that Amazon invests in new business opportunities and the timing of those investments, the mix of products and services sold to customers, competition, management of growth, international growth and expansion, the outcomes of claims, litigation, government investigations, and other proceedings, data center optimization, variability in demand, the degree to which we enter into, maintain, and develop commercial agreements, and proposed and completed acquisitions and strategic transactions. Other risks and uncertainties include, among others, risks related to new products, services, and technologies, system interruptions, and government regulation and taxation. In addition, global economic and geopolitical conditions and additional or unforeseen circumstances, developments, or events may give rise to or amplify many of these risks. More information about factors that potentially could affect Amazon's future business, product development, and financial results is included in Amazon's filings with the Securities and Exchange Commission, including its most recent Annual Report on Form 10-K and subsequent filings.

Top related stories What generative AI means for businesses and how AWS can help How generative AI helps Amazon eliminate checkout lines and revolutionize the shopping experience Why this AWS VP believes generative AI has the potential to transform our lives How generative AI helped train Amazon One to recognize

your palm

Related Tags

Share