

# Snowflake's Consumption-Based Cloud Model

## Idea In Short

Snowflake charges customers the way a utility company does: for what they actually use, measured down to the second, rather than for a fixed subscription. That consumption-based pricing model, combined with an architecture that separates data storage from computing power, let Snowflake grow from a stealth startup founded by three former Oracle engineers into a company that posted 4.68 billion dollars in revenue for its 2026 fiscal year. The company went public in 2020 in one of the largest software initial public offerings (IPOs) on record, then replaced its founding-era chief executive in 2024 to chase a new opportunity in artificial intelligence (AI). This article traces Snowflake's technical architecture, its pricing model, its record-setting public listing and its more recent pivot toward AI-driven products under new leadership.

-

## Three Oracle engineers bet on the cloud

Benoit Dageville, Thierry Cruanes and Marcin Zukowski founded Snowflake in San Mateo, California, in 2012, bringing with them years of experience building database systems at Oracle. Dageville had spent 15 years at Oracle as a lead architect for parallel execution, while Cruanes worked 13 years there on query optimization before the two decided existing data warehouse technology, built for on-premises servers, was poorly suited to the cloud. They named the company Snowflake partly as a nod to their shared enthusiasm for snow sports and partly because no two snowflakes look alike, a detail meant to reflect the tailored approach to data management they wanted to build.<sup>1</sup>

The company operated in stealth mode for its first two years, revealing itself publicly in October 2014 alongside a 26 million dollar funding round. That quiet development period let the founders build a working product before facing competitive scrutiny from the established database vendors they had just left.

## **An architecture that splits storage from compute**

Snowflake's central technical decision was to separate data storage from the computing resources used to query that data, a departure from traditional data warehouses that bundled the two together in a single fixed system. Customers create virtual warehouses, which function as independent clusters of computing power and can scale each one up or down without touching the underlying stored data. Because storage and compute scale independently, a company running a large overnight data-processing job does not need to pay for that same computing capacity the rest of the day when only small queries are running.<sup>2</sup>

That architecture also underpins Snowflake's multi-cluster, shared-data design, which lets many customers query the same underlying data simultaneously while keeping each customer's access isolated and secure. The design has proven flexible enough to support workloads far beyond the traditional data warehousing Snowflake started with, including the machine learning and AI applications the company now emphasizes.

## **Billing by the second instead of the year**

Snowflake launched its first commercial product, a cloud data warehouse, to the public in 2015 after raising 45 million dollars in funding that year. In 2016, the company introduced a usage-based pricing model that charged customers for the exact amount of computing and storage they consumed, billed in per-second increments, replacing the fixed annual subscription pricing that had defined the data warehousing industry until then. That shift mattered because it removed the tradeoff customers previously faced between overpaying for unused capacity and risking performance problems during unexpected demand spikes.

The pricing model also changed how Snowflake's own revenue behaves. Because customers pay based on actual usage rather than a fixed contract value, Snowflake's revenue tends to track how intensively its largest customers use the platform, which is part of why the company reports net revenue retention, a measure of how much existing customers increase their spending year over year, as closely as it reports new customer growth.

## **Multi-cloud partnerships opened new markets**

Snowflake partnered with Microsoft Azure in 2019, letting customers run Snowflake's data warehouse on Microsoft's cloud infrastructure rather than only on Amazon Web Services, which had hosted the platform since its founding. The company extended that multi-cloud strategy in 2020 with a partnership with Salesforce and again in 2021 with Google Cloud Platform, giving customers the option to run Snowflake on any of the three major public cloud providers. That same year, Snowflake introduced Snowpark, letting developers build data applications directly on the platform using Python, Java or Scala and Snowsight, a built-in data visualization tool.<sup>3</sup>

The multi-cloud approach reduced a risk that had worried some enterprise customers: dependency on a single cloud vendor for both infrastructure and data warehousing. By running on Amazon Web Services, Microsoft Azure and Google Cloud Platform simultaneously, Snowflake let customers choose the infrastructure they already had contracts with, rather than forcing a migration as a condition of adopting Snowflake's software.

## **A record-setting debut on Wall Street**

Snowflake priced its initial public offering (IPO) at 120 dollars a share in September 2020, above an already-raised expectation range and the stock opened trading at 245 dollars before closing its first day at 253.93 dollars, more than doubling its offer price. The debut valued Snowflake near 70 billion dollars and made it the largest software IPO in history at the time, a striking outcome for a company that was still recording significant losses. Investors were betting on the growth trajectory of consumption-based revenue rather than current profitability, a wager that played out through the stock's continued volatility over the following years.<sup>4</sup>

Snowflake's stock more than doubled on its first day of trading in September 2020, closing near 254 dollars against a 120 dollar offer price

The IPO gave Snowflake capital to accelerate product development and sales expansion, funding that supported the Snowpark and Snowsight launches the following year and the acquisitions that followed later in the decade.

## **New leadership, new pitch: the AI Data Cloud**

Frank Sloatman, who led Snowflake through its IPO and early public-company years, retired as chief executive on February 28, 2024 and the company named Sridhar Ramaswamy, who had joined Snowflake months earlier through its acquisition of the AI search startup Neeva, as his replacement. The stock fell more than 20 percent the day of the announcement, reflecting investor uncertainty about the leadership transition even as the company reported strong quarterly results alongside it. Ramaswamy, who had previously led Google's advertising products for 15 years, has since steered Snowflake toward a rebrand as an "AI Data Cloud", launching the open-weight Arctic language model, the Cortex AI suite of tools and Snowflake Intelligence, an agentic AI product that reached general availability in November 2025.<sup>5</sup>

The company backed that pivot with acquisitions, including Observe for roughly 1 billion dollars along with smaller deals for Select Star, Datometry and Natoma, each aimed at extending Snowflake's platform toward AI-ready data management rather than warehousing alone. For fiscal year 2026, Snowflake reported total revenue of 4.68 billion dollars, up 29 percent from the prior year, with net revenue retention of 125 percent, indicating existing customers continued expanding their usage substantially.

## **Competing against the cloud providers it depends on**

Snowflake competes directly with data warehouse products sold by the same cloud providers whose infrastructure it runs on, including Amazon Web Services' Redshift, Google's BigQuery and Microsoft's Azure Synapse Analytics. Oracle, the company all three Snowflake founders previously worked for, competes through its Autonomous Data Warehouse product, while Teradata, founded in 1979 and IBM's Cloud Pak for Data serve overlapping enterprise customers with longer operating histories in data management. Snowflake has differentiated itself by emphasizing simplicity and cross-cloud flexibility rather than being tied to a single provider's ecosystem, a positioning that matters most for large enterprises unwilling to commit their entire data infrastructure to one vendor.

That competitive dynamic creates a structural tension in Snowflake's business:

its growth depends partly on the willingness of Amazon, Google and Microsoft to keep hosting a product that competes with their own data warehouse offerings,

since Snowflake still runs on their infrastructure rather than owning its own data centers

## **Key Partners**

Cloud infrastructure providers, including Amazon Web Services, Microsoft Azure and Google Cloud Platform, host the servers and storage that Snowflake's platform runs on, making them essential partners despite also competing with Snowflake's own products. Technology partners such as Tableau, Looker and Databricks integrate with Snowflake to extend its analytics and visualization capabilities for joint customers. Consulting firms and system integrators help large enterprise clients implement and customize Snowflake deployments, particularly for complex migrations from legacy data warehouses. Value-added resellers extend Snowflake's distribution into markets and customer segments its direct sales team does not cover as efficiently.

## **Key Activities**

Snowflake's core activity is operating and continuously improving its cloud-based data platform, which requires ongoing investment in the underlying architecture that separates storage from compute. Research and development consumes a significant share of the company's resources, reflecting the pace of new features, from Snowpark to Cortex to Snowflake Intelligence, released in recent years. Sales and marketing activities focus on converting enterprise prospects, often already using a competing data warehouse, into Snowflake customers. Customer support and success operations help large accounts optimize their usage and manage costs under the consumption-based pricing model.

## **Key Resources**

Snowflake's proprietary architecture, which separates storage from compute and enables multi-cluster access to shared data, is the technical resource that differentiates it from older data warehouse products. Access to cloud infrastructure from Amazon Web Services, Microsoft Azure and Google Cloud Platform gives Snowflake the physical computing capacity to deliver its service without owning data centers. Its patents and other intellectual property protect the underlying technology from being directly copied by competitors. A

workforce of experienced data engineers, cloud infrastructure specialists and AI researchers, many recruited through acquisitions like Neeva, supports continued product development.

## **Value Propositions**

Snowflake offers customers a data warehouse that can scale storage and computing power independently, so businesses only pay for the capacity they actually use rather than a fixed system sized for peak demand. Its platform works across multiple data types and formats, including structured, semi-structured and unstructured data and integrates with common data pipeline tools, reducing the engineering work needed to consolidate information from different sources. Security and compliance features, including end-to-end encryption and adherence to standards like SOC 2, HIPAA and GDPR, address concerns that might otherwise block adoption among regulated industries. Its cross-cloud availability lets customers avoid being locked into a single cloud infrastructure provider for their data warehousing needs.

## **Customer Relationships**

Snowflake assigns dedicated account managers and customer success teams to larger accounts, supporting them from initial onboarding through ongoing troubleshooting and usage optimization. An active user community shares implementation practices and solutions to common problems, reducing the support burden on Snowflake's own staff for routine questions. The company gathers customer feedback systematically through surveys and account reviews, feeding it into product roadmap decisions. Technical support tiers vary by customer size and contract value, with the largest enterprise accounts receiving more direct engagement from Snowflake's engineering and product teams.

## **Channels**

Snowflake sells primarily through a direct sales team that engages enterprise prospects, particularly those already running data infrastructure that Snowflake's platform could replace or augment. Partnerships with cloud providers, software vendors and system integrators generate a meaningful share of new customer introductions, since many enterprises discover Snowflake through an existing vendor relationship. An online self-service channel lets smaller customers sign up, access documentation and manage their

accounts without direct sales involvement. Industry conferences and Snowflake's own annual Summit event serve as a channel for both customer acquisition and product announcements.

## **Customer Segments**

Large, mid-market and small businesses across industries make up Snowflake's primary customer segment, drawn by the platform's scalability and security for managing substantial data volumes. Government organizations, including agencies like the U.S. Department of Defense and the U.S. Census Bureau, use Snowflake for data management needs that require strict security and compliance guarantees. Non-profit organizations represent a smaller but ongoing customer segment, typically drawn by Snowflake's flexible pricing for organizations with variable data processing needs. Increasingly, customers building AI and machine learning applications form a distinct and growing segment, driven by Snowflake's Cortex AI suite and Snowflake Intelligence products.

## **Cost Structure**

Storage costs make up part of Snowflake's cost base, reflecting the cloud infrastructure fees it pays to host customer data across its partner providers. Compute costs vary with usage, since Snowflake's own infrastructure expenses rise and fall with the volume of customer queries processed on the platform at any given time. Research and development spending remains a significant and growing cost category, given the pace of new AI-related product launches the company has pursued since 2024. Sales, marketing and customer support costs round out the structure, reflecting the enterprise-focused sales motion Snowflake relies on to win and retain large accounts.

## **Revenue Streams**

Cloud data platform subscriptions, billed on measured usage of storage, compute and data transfer, generate the large majority of Snowflake's revenue. Professional services, including training and consulting for customers implementing or optimizing their Snowflake deployments, add a smaller supplementary revenue stream. The Snowflake Marketplace generates commission-based revenue by connecting customers with third-party data and application providers. Snowpark and partnership-driven sales, where Snowflake earns a commission on transactions generated through its technology and reseller partners, round

out the company's remaining revenue streams.

- 1Snowflake co-founder Benoit Dageville's Oracle background
- 2Snowflake's early stealth-mode funding and product strategy
- 3Snowflake's expansion across cloud infrastructure providers
- 4Snowflake's 2020 IPO and first-day trading performance
- 5Snowflake's 2024 CEO transition and market reaction

## Summary

Snowflake's business model succeeds because it removed a tradeoff customers used to accept: pay for fixed capacity you might not need, or risk running out of capacity during peak demand. By separating storage from computing and billing per second of use, Snowflake let customers scale each independently and pay only for what their workloads actually consume. That architecture, more than any single feature, is what convinced enterprises, government agencies and cloud providers themselves to build on the platform. The leadership change in 2024 and the subsequent push into artificial intelligence (AI) products like Cortex and Snowflake Intelligence mark an attempt to extend that model beyond storage and query processing into a broader data and AI platform. Competition from Amazon Web Services, Google and Microsoft, all of which sell competing data warehouse products alongside the cloud infrastructure Snowflake depends on, remains the clearest risk to that strategy.