

AdAge

UNLOCKING THE POWER OF THE DATA CLOUD

Sponsor Content



UNLOCKING THE POWER OF THE DATA CLOUD

Marketers are increasingly turning to cloud-based technology platforms that can leverage vast sums of customer data into real-time insights for optimizing personalized content, advertising and digital media strategy.



iStock

Today's marketers are at the forefront of an evolving conversation around data. By now, most successful companies recognize the value of big data. Many have taken steps to expand their data portfolios and strengthen their analytics capabilities. Increasingly, the challenge they face is how to source, store, analyze and share vast sums of customer data from multiple sources across the organization in a way that produces insights in real time to inform critical business decisions.

With the right systems and processes in place, data can help deliver the growth and innovation every marketer from Fortune 500 companies to early-stage startups need in order to compete in today's economy. Data is no longer one asset among many: It is the key to gaining a competitive edge in customer

relationships, delivering engaging and personalized customer experiences, anticipating new market trends, predicting consumer behavior and developing new business strategies—to name just a few critical areas. Everything from personalized content to advertising and digital media strategy can be optimized with data.

The need for advanced data capabilities spans across many different industries and marketing applications. Media, entertainment and gaming companies, for example, are seeking to capitalize on the explosion of data created by streaming and OTT services, whose platforms allow marketers to build direct relationships with consumers through one-to-one ad targeting and personalized experiences. Companies from technology to retail/CPG and financial services are all looking for ways to engage more

**The data cloud:
A global network**

The data cloud allows organizations to source and derive value from rapidly growing data sets regardless of where that data is stored.



Courtesy Snowflake

deeply with consumers using new digital technologies and platforms fueled by data.

Despite these powerful market incentives, however, many companies struggle to turn data into insights. Further research reveals a lingering trust gap over the reliability of data. For example, in a January 2019 survey of 2,400 global executives by MIT, more than three-quarters of respondents reported increased access to data over the previous two years, yet only 43 percent said they have the “right” data—i.e., the data that drives decision-making. Moreover, only a small fraction said their data is always accurate (9 percent) or complete (6 percent).

So how can companies begin to close this gap? In short, having the right data and the ability to leverage it effectively depend on having right data architecture. And that foundation starts where so many organizations are looking to bring their teams into the future: the cloud.

Understanding the shift to the cloud

Many factors underpin the corporate world’s shift to the cloud. The rise of cloud-based workloads has led to a proliferation of connected devices, applications and social media, resulting in an explosion of digital data. According to IDC, there will be 175 zettabytes of data—a zettabyte is a unit of storage capacity that represents by 10²¹, or a trillion gigabytes—by 2025, representing a compound annual growth rate of 27 percent from 33 zettabytes of data in 2018. Nearly half (49 percent) of all data will be stored in public

cloud environments by 2025, up from about 30 percent today.

Conventional data warehouses (i.e., any computer system dedicated to storing and analyzing data to reveal trends, patterns and correlations that provide information and insights) were not designed to handle the volume, variety or velocity of today’s available data. Traditional on-premises data warehouses from Oracle, Teradata and Netezza, for example, can be rigid and difficult to use and limited in computation and storage, depending on a company’s available hardware and physical server space.

These legacy systems present a number of operational challenges. Most are fixed, meaning that each data “box” has a certain amount of built-in storage and computing capacity, often forcing a company to expand its infrastructure and buy into additional capacity that it may never use. Those resource requirements can change on a daily, if not hourly, basis. As an example, consider the needs of a major chain retailer as it integrates all of its customer data to support cross-functional teams on Black Friday. In that case, the company could conceivably end up paying for computing and storage needs on a single day what it actually needs for the rest of the year. And it isn’t just a matter of cost, either. Legacy systems have given rise to data silos, which, like organizational silos, further prevent the optimal sharing and application of data (see “Breaking Down the Data Silo,” page 5).

Fortunately, there is a better solution. It lies within the data cloud.

Snowflake data cloud platform

SaaS architecture enables organizations to work on common data for a wide range of workloads.

Data sources

- OLTP databases
- Enterprise applications
- Third-party
- Web/log data
- IoT



Data consumers

- Data monetization
- Operational reporting
- Ad hoc analysis
- Real-time analytics

Courtesy Snowflake

The data cloud is a global network of data that sits at the intersection of the infrastructure cloud (e.g., services provided by Microsoft, Amazon, Google) and the application cloud (Salesforce, Workday, ServiceNow). The data cloud renders individual silos obsolete by combining disparate sources of both structured and semi-structured data into a single location, according to Matt Feldman, VP of enterprise sales at Snowflake. “If we can connect customer data and supplier data with third-party providers, as well as data partners, then we can potentially unlock a tremendous amount of business value,” he says.

In recent years, companies like Snowflake have emerged to become the engine behind the data cloud. Unlike legacy databases or on-premises data warehouses, Snowflake is software-as-a-service (SaaS), so instead of having to maintain physical hardware and equipment, or manually create local updates to the most recent MySQL server (a commonly used open-source relational database management system), Snowflake updates its cloud data warehouse through a web user interface. The Snowflake cloud data platform is built specifically for the cloud and leverages the speed, scale and flexibility that the cloud permits. It is the one place where companies can pipe in their application and SaaS data to analyze and make business decisions quickly and reliably, using SQL or a different reporting tool like Tableau.

“We can do this not only across multiple cloud platforms—including Amazon’s AWS,

Google’s GCP and Microsoft Azure—but also multiple regions within those clouds,” notes Feldman. “We’re giving companies access to data from across the globe.”

Snowflake’s growth is reflected in the expanded overall market opportunity for cloud data warehouse solutions.

Privacy and data sharing trends accelerate adoption

Recent trends in online advertising are accelerating the move to cloud-based platforms. Facing new restrictions on data mining from third-party sources, such as the removal of cookies from Google’s Chrome browser and more stringent government-enforced privacy regulations, advertisers are seeking to gain broader access to first-party data to enrich their data sets in a privacy-compliant way. As a result, there is growing demand for “matching” first-party data between publishers and advertisers within newly created private data exchanges and global data marketplaces, according to Bill Stratton, head of media, entertainment and advertising at Snowflake.

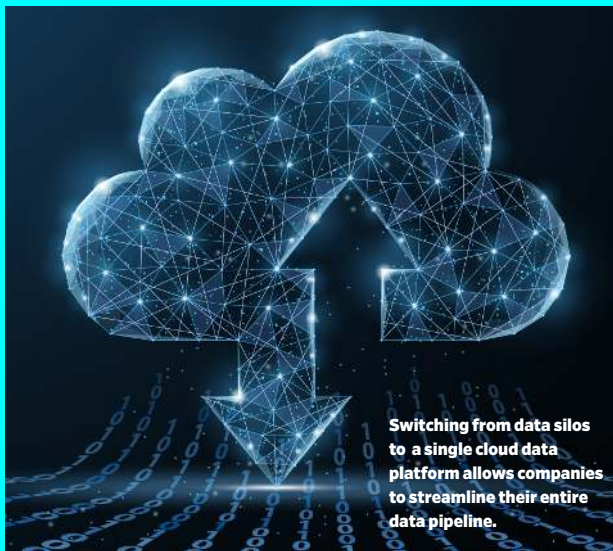
“The Snowflake data marketplace is very different—and more secure—than cookie data marketplaces,” says Stratton. “The data is shared with the customer from the provider without the physical data itself having to be picked up, replicated or moved. Instead of using cookies, the matching takes place on identifiers that the ad partners share in common, such as a hashed email or mobile ad ID.”

BREAKING DOWN THE DATA SILO

What exactly are data silos, and why are they a problem for marketers? In short, data flows into an organization, is stored in structures such as data marts and data lakes—two common types of silos—and is then activated in reporting and analytics as well as in driving customer engagement for things like personalized offers, e-mail and advertising. For these efforts to succeed, product, sales and marketing teams all need to be working off the same sets of data. But in a world where data is stuck in organizational, functional or technical silos, this collaboration cannot occur.

Let's say a media company wants to send an ad to motivate more usage of Hulu's mobile app and needs to know whether users are fans of documentaries or animated films. This type of content consumption data often resides in the product department. But to drive personalization, it needs to be available to the marketer at a granular level (not aggregates) and in real-time (not copied and moved).

For a media and technology company focused on life milestone events, siloed data inhibited the development of dimensional models that could deliver meaningful insights to internal stakeholders and advertisers. The company needed to simplify the collection and analysis of semi-structured data, including page views, clicks and interactions from web and mobile apps. Its on-premises warehouse could only handle structured data, such as membership and



Switching from data silos to a single cloud data platform allows companies to streamline their entire data pipeline.

iStock

vendor information, so the team needed a way to convert semi-structured interaction events into actionable analytics.

Initially, the team experimented with a second data warehouse. “Maintaining a dedicated platform for semi-structured data wasn’t scalable or cost effective,” says an executive at the company. Snowflake’s multi-cluster cloud data architecture aligned with the company’s rapidly expanding data needs, and its separate storage and compute capabilities eliminated the potential for future resource contention.

Ingesting everything into a single cloud data platform democratized data analytics at the company by allowing the team to rapidly build and iterate dimensional models that feed business intelligence dashboards. Democratized data enabled product development and marketing to measure the effectiveness of product features, promotional campaigns and retention efforts. Advertisers were able to receive timely performance statistics about storefront visibility, interaction rates and comparative advertising.

Breaking down the data silos was a clear breakthrough for the company. Switching to the cloud “streamlined our entire data pipeline,” says the company executive.

Today's advertisers are also seeking safe-haven mechanisms like clean rooms, which allow brands to access data from the walled gardens of Facebook, Amazon and Google without sharing the information outside the two companies, and to buy and sell ads more frequently. "Snowflake is not an ad tech company. We don't activate campaigns. We integrate with the many companies that do," explains Stratton. "Once the ads are up and running, the marketer needs a place to bring all of the campaign data back to analyze. We provide that environment. The more devices and platforms proliferate, the more data from each source is brought back into the Snowflake ecosystem."

There are many potential benefits that accrue to an organization when switching from a legacy data warehouse to a cloud data system. The top three are scalability, speed/performance and concurrency.

• Scalability

Cloud data platforms are built on top of the public cloud infrastructure, which enables near limitless scale. This differs from traditional on-premises data warehouses, which have physical limitations based on actual physical server size. For one large hospitality and entertainment company, the transition from its old data centers to the Snowflake Cloud Platform freed up time and resources and allowed the company to focus on other critical areas of the business, which includes growing its catalog licensing revenue.

"We evaluated a lot of providers and found that Snowflake was quite labor-less in its scalability," says a senior executive at the company. "Throughout the year we do what we call 'avails' calculations—what's available to sell in what territory and when. Some of these calculations are quite intense when you look at the size of the catalog we have. Snowflake handled those with ease."

The company also looked to Snowflake to help pull in new data sets and scale those up for insights leveraged by marketing and sales teams. As the executive explains, "We said, let's pull in all the feeds we have on buy data and social sentiment and make decisions coming out of our data warehouse on what titles

do well on which platforms. If we're green-lighting a new project, what title should it go on? Should it be exclusive or should we spread it over multiple platforms?"

• Speed/performance

At a fundamental level, the cloud can offer a lot more speed or performance, thanks to the sheer volume of computing resources the company can throw at any given job. "Before it would take us several days to do a global avails run. Now it takes a few hours," the entertainment executive notes. "That speed allows us to pull data back for sales and sell an existing catalogue more efficiently."

Electronic Arts is one of the largest video game publishers and developers in the world. With business transitioning from the traditional retail to digital downloads, subscriptions and mobile games, the volume of data and data velocity has grown exponentially. New revenue models require the addition of player activity metrics into the same analytical space—multiplying the pressure on the platform. The existing legacy IT data platform of the past was not able to scale to the new demand. Snowflake offered a cloud-based SaaS data platform that could easily handle the demand for the volume and speed. The queries that ran for multiple minutes and in some cases hours on the legacy platform came back in seconds and even milliseconds in the cloud. Snowflake also allowed for the unique ability to scale storage and compute independently based on the actual demand and scale it back when the demand diminishes seasonally, paying only for what has been used.

• Concurrency

Concurrency is the ability to perform two or more tasks simultaneously, or to allow two or more users access to a computing solution. Cloud-optimized architectures support concurrency in the following ways: (1) Multiple users can query the same data simultaneously without degrading performance. (2) Loading and querying can happen concurrently, thereby enabling and supporting multiple, simultaneous workloads without resource contention.

Companies like Asics are constantly running data pipeline jobs in the background moving data from, say, Google and Facebook into their

data warehouse. A problem they previously ran into was they couldn't concurrently run data cleaning and transformation scripts and query their data at the same time (how frustrating!). Snowflake allowed ASICS to do that by having a separate warehouse so the analysts don't have to wait for the data engineers ETL jobs to finish before running their report.

What to look for in cloud partners—and the future

As every marketing manager knows, today's consumer is a constantly changing target and promotional success can turn on a dime. Event-driven analytics demand constant data so that managers can update reports and dashboards on a continual basis and monitor performance in real time—a process best handled by an elastic cloud data warehouse that can accommodate variations and spikes in data flow.

Not all cloud data warehouses feature the same degree or type of elasticity. Advanced solutions can scale up and down, on the fly and without taking the system offline or putting it into a read-only mode. Conversely, a cloud data warehouse that requires manual reconfiguration involves careful planning and coordination with the vendor to scale resources. Scaling may require downtime or a switch to read-only mode to redistribute data and reconfigure the system. Most cloud data warehouse offerings bundle computing and storage on the same node, requiring customers to scale both when they need to increase just one or the other. In reality, many of these offerings are simply "cloud-washed" versions of on-premises solutions that are not designed for peak usage and therefore will likely demand costly upgrades.

Elasticity is one key factor in selecting a cloud data technology partner, but there are many others. When choosing solutions built on a cloud-optimized architecture, companies should look for the following characteristics:

- Centralized storage for all data
- Independent scaling of compute and storage resources
- Near-unlimited concurrency without competing for resources
- Loading and querying data simultaneously without degrading performance
- Replicating data across multiple regions and clouds to enhance business continuity and simplify expansion
- Sharing data without setting up APIs or establishing cumbersome ETL procedures

As with any digital transformation, moving out of legacy data warehouses and into the cloud requires buy-in from stakeholders across the organization. Any ROI assessment should be framed in terms of reducing costs, improving efficiencies and driving growth. Switching to the cloud allows for a per-second billing model in which companies only pay for what they use. It saves time in many critical data functions—including managing data pipelines, governance, security, backup and disaster recovery—and enables teams across the company to access data in order to drive revenue and profit.

Throughout the industry, the appetite for data has never been bigger. And now, thanks to the data cloud, connecting data and sharing it across the organization is no longer the daunting task it used to be. The traditional barriers to effective data management—moving data, copying data, data sharing and data privacy—cannot be solved by old technology. The paradigm is changing. The data cloud allows a single copy of data to be virtualized and applied in near infinite use cases for just about any business or marketing objective. The organization's incentive for adoption is not just a matter of creating efficiencies in the short term, either. Companies have only begun to realize the potential long-term impacts on growth and innovation that can result from transitioning their data-sharing practices to a modern cloud-based platform. For these brave companies, quite literally, the sky's the limit.

Ad Age Studio 30

Ad Age Studio 30 helps your brand connect with an influential audience actively seeking new partners, solutions and products. Through original custom articles, thought-leadership content, events, research, webcasts, white papers, infographics and more, our end-to-end solutions help your content reach and resonate.

Studio30@adage.com

Staff:

Writer: **Michael Applebaum**

Senior Art Director: **Jennifer Chiu**

Copy Editor: **Brian Moran**

Contact us:

James Palma

General Manager, Revenue

and Client Partnerships

jpalma@adage.com

John Dioso

Editor, Studio 30

jdioso@adage.com

About Snowflake

Snowflake delivers the Data Cloud — a global network where thousands of organizations mobilize data with near-unlimited scale, concurrency and performance. Inside the Data Cloud, organizations unite their siloed data, easily discover and securely share governed data, and execute diverse analytic workloads. Wherever data or users live, Snowflake delivers a single and seamless experience across multiple public clouds. Join Snowflake customers, partners and data providers already taking their businesses to new frontiers in the Data Cloud.

www.snowflake.com