

What Is Zero Copy? (and How It Works)

How can a customer data platform complement your data warehouse? By providing instant access to data without a lift-and-shift.



Martin Kihn

April 24, 2024 · 8 min read



Zero copy integration makes it possible to access data that is sitting in multiple different databases at the same time without having to move, copy, or reformat anything. [Getty Images / Studio Science]

Imagine if your furniture and belongings could just teleport to your new place in perfect condition. It's not possible (yet) in the physical world, but with [zero copy integration](#), that's how you can [handle your customer data](#).



Thanks to zero copy or zero ETL (extract-transform-load), it's possible to share data among two or more data stores without actually moving it. This is great news to companies that store data in a cloud data warehouse like Snowflake or Google BigQuery. Some of them are reluctant to adopt a [customer data platform](#) (CDP) because they don't want to duplicate data.

They don't have to. Using zero copy integration, users can get the [benefits of a CDP](#) – like data harmonization, identity management, built-in analytics and activation – without the downside of physical data movement.

What you'll learn

- [What is zero copy integration?](#)
- [What is a data warehouse?](#)
- [How it works: from CDP to data warehouse](#)
- [How it works: from data warehouse to CDP](#)
- [How Buyers Edge uses zero copy technology](#)

What is zero copy integration?

Zero copy integration makes it possible to access data that is sitting in multiple different databases at the same time without having to move, copy, or reformat anything. In addition to making access faster and easier, it cuts down on the expense and risk of errors that's always incurred when data has to be moved or changed.

Copying data from one database to another is a common practice. Often, this process entails some form of data transformation called extract-transform-load (ETL). It can be a useful and even necessary step in [managing enterprise data](#).

But it has its challenges. Some of the differences between traditional (copying) methods and the zero copy approach are:

	Traditional	Zero Copy
Replication	Source data copied from original location to target	Data remains in original location
Updates	Data only accurate as of last synchronization point	Data is accessed in real-time
Cost	User pays cost of moving and syncing data	No data movement cost
Regulatory requirements	Harder to keep up with compliance due to more complex governance	User only responsible for source data
Errors	Any data movement introduces potential for errors or mistakes	No movement errors
Maintenance	Copying and syncing creates more complexity	Easier to manage

Typically, the physical copying of data incurs costs for data transportation, introduces the potential for errors, complicates data governance and management, and creates data-syncing time lags.

So how does zero copy integration work? The actual mechanism differs from platform to platform and is different whether you are accessing data from the CDP into the data warehouse or vice versa.

In the following examples, we'll be using Salesforce Data Cloud as the CDP and [our partner](#) Snowflake as the data warehouse. Other vendors could be substituted without significantly

Just For You



[What Does Salesforce Do?](#)

11 min read



[3 Ways to Use B2B Data to Simplify Your Marketing Campaigns](#)

6 min read



changing the explanation.

[\(Back to top\)](#)

Dig deeper



How Data, Trust, and Contextual UX Deliver AI Success



WATCH: How To Build Targeted Advertising In 3 Steps



A CDP Is a Powerful Tool — Here's How To Make It Work for You



9 Ways AI Can Save Marketers Time, Money — and Grief

What is a data warehouse?

A data warehouse is simply a reliable place to [store and access data](#) that is important to the business.

Traditional data warehouses work with highly-structured data in formatted tables, and they tend to be quite slow and complicated. On the other hand, modern data warehouses like Snowflake can handle almost any type of data, process it quickly, and are easier to use. Because they are built on top of cloud platforms like Amazon and Google, they are easier to plug into other systems like CDPs that use the same platforms.

[\(Back to top\)](#)

How it works: from CDP to data warehouse

In this case, we are inside our data warehouse and want to access data that is in the CDP. In other words, information is going *out* from the CDP to the data warehouse. This process is sometimes called [data sharing](#).

The usual steps are:

1. Identify the objects – or data nuggets – within the CDP you'd like to share. In the case of Salesforce Data Cloud, these are called data lake objects (cleansed data), data model objects (structured by the CDP user for their business cases), and calculated insights objects (for formulas like lifetime value).
2. Using point-and-click, link these objects to the data share target, in this case Snowflake.
3. Inside Snowflake, the user can perform queries across data in Snowflake as well as the objects linked via the data share – all at the same time.

Behind the scenes, the process creates “virtual tables” that describe the Data Cloud data to Snowflake. A virtual table is like a window into data in a database, but instead of copying and storing actual data, a virtual table only contains the *structure* of the data. It's a blueprint or pointer to the right place in the CDP to get the data – but the data itself stays in the CDP.

“It is possible to query live data in Salesforce from Snowflake and ensure that changes in the Salesforce objects will be reflected in Snowflake,” explained Salesforce Data Cloud product manager Sriram Sethuraman. “This will empower developers and data scientists to build machine learning models and [AI-powered applications](#) on top of the Snowflake platform by joining Salesforce and Snowflake data.”

[\(Back to top\)](#)

How it works: from data warehouse to CDP

Now we are inside our CDP and would like to access data that is sitting in our data warehouse. This process is sometimes referred to as [data federation](#).

There are a lot of good reasons to do this. Data warehouses like Snowflake and Google BigQuery usually contain a massive amount of data, including transactional data like purchases, and product data. Although not typical “customer” data, such information can be very useful when trying to calculate a customer's loyalty status or build a recommendation based on details about products they buy.

For example, here's how you can access data warehouse data [in Salesforce Data Cloud](#):

1. Salesforce Data Cloud mounts tables from the data warehouse as external data lake objects. (Mounting is a process that creates a virtual data blueprint, like the one described above.)
2. Data Cloud performs its usual functions such as ID management, analysis, segmentation, etc.

3. The CDP can access data from the data warehouse by performing federated (or combined) queries that include data in Data Cloud and the objects that are provided by the data warehouse.

[\(Back to top\)](#)



How Buyers Edge uses zero copy technology

The success of cloud-native data warehouses like Snowflake, Databricks, Google BigQuery and Amazon Redshift makes a lot of sense. We've seen many customers at least experiment with them and many use them as an integral part of their data architectures. But no data warehouse performs all the functions of a CDP, such as identity management and user-friendly analytics.

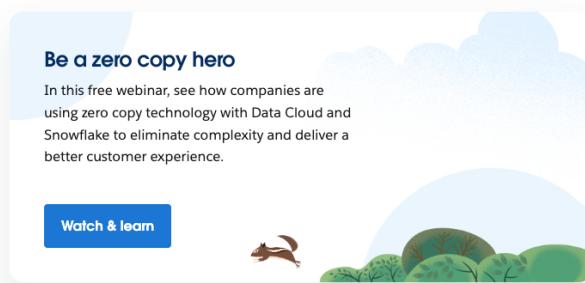
[Buyers Edge](#) – a leading procurement optimization company in the food service industry – wanted to build a unified customer profile in a CDP while accessing purchase data stored in a data warehouse. Their main goal is to provide better customer insights back to their sales and marketing teams.

Using the zero-copy connection between Data Cloud and their warehouse, Buyers Edge gains access to the purchase data it needs to build predictive models, allowing sales and marketing teams to produce better offers, messages and experiences for its prospects and customers.

"With zero-copy technology, accessing customer data stored in Salesforce becomes effortless, eliminating the need for data movement, duplication or reformatting," said Sean Donahue, chief of staff for the Buyers Edge Platform. "This saves time and resources and removes data silos, harmonizes data for insights and analytics, and empowers businesses with a real-time holistic view of our customers."

And as companies like Buyers Edge evolve, their requirements will change. That's why a technology like zero-copy can help them and others build a more flexible data management strategy.

After all, larger enterprises have an average of [976 different applications](#) running their business, and the amount of data created, captured, copied, and consumed is expected to more than [double by 2026](#). Thanks to the power of zero copy data sharing, the looming data explosion will be a lot easier to enjoy.



[\(Back to top\)](#)

Share article



Explore related content by topic





Martin Kihn
SVP, Market Strategy, Marketing Cloud

Martin Kihn is the Senior Vice President of Market Strategy for Marketing Cloud. In a former life, he was a research vice president at Gartner, where he wrote and spoke widely about marketing technology, and advised numerous Fortune 500 clients on marketing strategy. He's also authored four books... [Read More](#)

[More by Martin →](#)



Get the latest articles in your inbox.

[Sign up now](#)

Just For You



[Marketers: What Not to Miss at Dreamforce '24](#)

3 min read



[Build Your 2024 Dreamforce Agenda](#)

4 min read



[Cyber Week Impact on Email Delivery: What Marketers Need to Know Now](#)

9 min read



[Why the Google Cookies News Shouldn't Alter Your Data Strategy](#)

7 min read

[Show More](#)



CALL US AT 1-800-664-9073

New to Salesforce?

- [What is Salesforce?](#)
- [Best CRM software](#)
- [Explore all products](#)
- [What is cloud computing](#)
- [Customer success](#)
- [Product pricing](#)

About Salesforce

- [Our story](#)
- [Press](#)
- [Blog](#)
- [Careers](#)
- [Trust](#)
- [Salesforce.org](#)
- [Sustainability](#)
- [Investors](#)
- [Legal](#)

Popular Links

- [Salesforce Mobile](#)
- [AppExchange](#)
- [Dreamforce](#)
- [CRM software](#)
- [Salesforce LIVE](#)
- [Salesforce for startups](#)



© Copyright 2024 Salesforce, Inc. All rights reserved. Various trademarks held by their respective owners. Salesforce, Inc. Salesforce Tower, 415 Mission Street, 3rd Floor, San Francisco, CA 94105, United States

[Legal](#) [Terms of Service](#) [Privacy Information](#) [Responsible Disclosure](#) [Trust](#) [Contact](#) [Cookie Preferences](#) [Your Privacy Choices](#)