



Solving ecommerce analytics

How retailers use data to understand customers, predict demand and optimize supply chains




Table of contents

00	Introduction: Data engineering & ecommerce analytics	3
01	Sleeping Duck Mattresses	4
02	carwow Cars	6
03	Westwing Home furnishings	8
04	Papier Greeting cards, stationary	10
05	Pet Circle Pet supplies	12
06	Koh Cleaning supplies	14
07	Ritual Vitamins, wellness products	16

INTRODUCTION

Data engineering & ecommerce analytics

Growing ecommerce businesses share a broad set of analytics challenges, regardless of product or industry. They need to understand customer behavior, predict demand, run highly targeted marketing campaigns and anticipate supply chain issues, to name a few.

Many of these challenges are existential — ecommerce companies can't compete successfully without solving them. They're also technically difficult. To create data models that illustrate the full customer lifecycle, for example, businesses need to join records across many disparate data sources. Smaller companies might not have the data engineering bandwidth for this kind of project — or might want to invest their engineering resources more strategically.

In this ebook, we take a close look at how growing ecommerce businesses in different industries are using a new data architecture — the modern data stack — to affordably solve this data engineering problem, and gain a significant competitive edge in the process.

A [modern data stack](#) (MDS) consists of automated data integration; a scalable, cloud-based destination like a data warehouse or data lake; and a modern business intelligence or data visualization tool. You can [set up an MDS in under an hour](#) and pay as little as a few hundred dollars a month for it. The technologies scale infinitely, so you won't have to continually engineer your analytics stack as you grow.

And as the following stories illustrate, an MDS can be invaluable for growth-minded ecommerce companies.



DISCOVER NEW MARKETS | STREAMLINE THE BUYER JOURNEY | OPTIMIZE AD CAMPAIGNS



S L E E P I N G D U C K

DATA STACK



Sleeping Duck

From its founding in 2014, Sleeping Duck has sought to gain a competitive edge in the crowded ecommerce mattress market by using data to guide its business decisions. But harnessing the power of its data proved to be challenging. Within a couple years, Sleeping Duck had data in many disparate sources: SaaS platforms, web apps, its own product, marketing platforms like Google Ads.

To gain a competitive edge via analytics, the company would need to pull data from all of these sources and analyze it holistically — a complex, manual process. Sleeping Duck didn't have the requisite data engineering resources in house.

"Our Fivetran data pipeline touches every department across the company. We are truly a data-driven company that can take advantage of ecommerce opportunities as they arise."

– Chao Li, data scientist,
Sleeping Duck

Data Scientist Chao Li chose Fivetran to centralize the company's data so he didn't have to worry about coding, configuration or ongoing maintenance of data pipelines. Fivetran could replicate analysis-ready data from dozens of sources into Snowflake, and included an ERD to summarize the data model for each SaaS source. This allowed Sleeping Duck to stay ahead of the competition with analytics projects like:

- **Optimizing marketing spend:** Going beyond counting clicks, the data team explores how leads turn into transactions and revenue. They know who is buying mattresses, what they're looking for, which pain points they're experiencing and why they choose Sleeping Duck over competitors.
- **Streamlining the customer journey:** The team can follow every customer through each purchasing phase — from first touch to mattress delivery — identifying bottlenecks and reshaping the buying journey so shoppers are ushered through the process to completion.
- **Improving customer support:** Service reps have a 360-degree view of customers who call with an issue, so they can drive brand loyalty by understanding whether customers have other open tickets and avoiding solutions that have failed in the past.



UNDERSTAND CUSTOMER BEHAVIOR | PREDICT FUTURE PERFORMANCE

carwow

DATA STACK



ECOMMERCE SUCCESS STORY #2

carwow

As a whole, the automotive industry lacks data — manufacturers build cars with very little consideration for what customers want. As a car sales platform and ecommerce innovator, carwow was determined to change that — it wanted to understand its customers and adapt to their behavior better than its competitors. To do so, however, it needed modern data architecture.

"Every morning when you come into work, yesterday's data is in Snowflake. Before Fivetran, you couldn't run queries on the previous day's data because the sync wasn't finished. That doesn't happen anymore."

– **Isaac Seymour,**
Senior Developer, carwow

Before its current data stack, different carwow teams performed DIY data integration for different data sources, time-consuming work that led to inefficient, unreliable data syncs. Data pipelines frequently broke, impacting multiple arms of the business, and rapid growth aggravated the problems.

With a modern data stack of Fivetran, Snowflake and Tableau, carwow can effortlessly centralize data from more than a half dozen sources into Snowflake. It no longer needs to build or maintain data pipelines, and SaaS data arrives ready for analysis, ERD included.

Carwow's data team uses this combined data to generate daily insights into buying behavior, customer preferences and conversions. Decision-makers now share real-time measurements of forward-looking performance — and carwow keeps pushing the analytics envelope in the automotive ecommerce space.



OPTIMIZE SUPPLY CHAINS | ENRICH CUSTOMER EXPERIENCE | IMPROVE MARKETING ROI

WESTWING

DATA STACK



Westwing

Ecommerce home furnishing provider
Westwing adopted a modern data stack
because it needed better insight into:

- Customer interactions with its website and product catalog
- Marketing spend ROI and the relative effectiveness of online channels
- The efficiency of key supplier and manufacturer relationships

"Now our investment in marketing is paid back several times because once we acquire the customer we can analyze everything they do and take very good care of them."

– **Christian Manrich, Business Intelligence Engineering Manager, Westwing**

Westwing had a relatively large in-house data team but still struggled to build and maintain data pipelines for every one of its many data sources — and those sources were only increasing as the business grew.

After adopting a modern data stack of Snowflake, Fivetran and Tableau, Westwing saved 40 data engineering hours a week while eliminating pipeline maintenance, improving data reliability and decreasing sync frequencies for key sources like Salesforce, Braze, Google Ads, Facebook Ad Insights and Google Sheets.

Westwing now has deep visibility into organic social media activity and paid marketing campaign performance, allowing it to optimize its online channels and increase ROI on paid campaigns. It analyzes CRM data from Braze and Salesforce to understand how to better streamline and nurture relationships with suppliers and manufacturers. And it can see how customers are interacting with its site and product catalog, enabling it to expand and improve the customer experience — critical for a company that depends heavily on customer loyalty.



CALCULATE CUSTOMER LIFETIME VALUE | BUILD A COMPLEX ATTRIBUTION MODEL



DATA STACK



Papier

Papier is a design and personalization business selling stationery, invitations, cards and photo books. Like many ecommerce businesses, it uses paid marketing as part of its acquisition strategy, so its ability to analyze spend across Facebook, Google, Bing and other ad providers is critical to its success. Before leveraging a modern data stack, however, Papier was trying to centralize ad, clickstream and transactional data into Redshift using ETL scripts and code written in house.

"A modern data stack has had such a transformational effect on how we look at data that it's almost difficult to imagine what we were doing before — we're so much more data-focused."

– Joe Robertson, CTO, Papier

That solution was far from optimal — the business was spending too much time on ETL chores, and CTO Joe Robertson and his team often found inconsistencies and inaccuracies while syncing historical data and matching data across campaigns.

A modern data stack of Fivetran, Redshift and Looker has allowed Papier to double the number of usable data sources its data team looks at regularly, and they can combine key transaction tables straight from its production database into its data warehouse.

This allows the data team to cross-reference transactional data with clickstream and event data and create a single source of truth for metrics on product views, orders placed and more. His team is able to calculate customer lifetime value by tying the transactional data together with site behavior data and other relevant data.

The Papier team has also used its modern data stack to combine data from ad providers and build its own attribution model. Its customer acquisition strategy is far more effective as a result.



MAP THE CUSTOMER JOURNEY | IMPROVE RETENTION | CREATE A SINGLE SOURCE OF TRUTH



DATA STACK



Pet Circle

Pet Circle, a major pet supply ecommerce business based in Australia, was looking to gain a competitive edge in a crowded market through data insights — specifically, it wanted to better understand the customer journey, dramatically improve the customer experience and create a customer-centric culture internally.

PetCircle hosted a customized stack in Google Cloud to support its subscription service model, but it was not centralizing its transactional data into a central warehouse.

"Having an analytical model to help us understand the direct impact of our decisions on customers is critical. A more sophisticated, customer-centric view will permeate the entire business. This is a game-changer for how we can impact our customers in a positive way."

– James Edwards, Head of Data Science, Pet Circle

Head of Data Science James Edwards chose Fivetran to replicate its MariaDB data into BigQuery, creating a single source of truth for analytics projects. It has confidence in the data and its data team never has to troubleshoot broken data pipelines.

The new analytics stack allows PetCircle to join views of website sessions with paid data in a single transformed table to see how often customers visit its site, log in, modify their subscriptions and more. That data helps Pet Circle improve loyalty and retention by adding customer value rather than transactional value.

Another project focuses on order metrics. Transactional systems don't deliver a single record of gross margin, so different teams throughout the business use SQL snippets to get their own numbers for gross margin. Having a single source of truth and reliable analytics eliminates data accuracy issues and enables the company to quickly understand the impact of its decisions on customers — a "game-changer," Edwards calls it.



IMPROVE CUSTOMER PROFILES | RUN TARGETED CAMPAIGNS | PREDICT CUSTOMER DEMAND

koh

DATA STACK



Koh

Cleaning supply ecommerce business Koh was expanding operations — without a clear data strategy. The data team was trying to generate insights by exporting data into spreadsheets and joining it there, a process that could take days. Koh lacked insight into customer behavior and struggled to run targeted campaigns, and it didn't fully understand the costs behind acquisition and revenue generation.

Increased customer
return rate from

25% to **50%**

So Koh adopted a modern data stack, using Fivetran to sync data from sources like Google Ads, Xero and Shopify into Google BigQuery and tapping Tableau for data visualizations and business intelligence.

Without having to code, configure or maintain data pipelines, Koh can join multiple data sets and create more complete customer profiles. By connecting purchase data from Shopify to data from Google Analytics, for example, it can see how customers arrive at its site and how many sessions it takes to convert them.

New customer insights allow the company to run highly targeted campaigns — by reaching customers on the right channel at the right time, it has doubled its customer return rate, from 25 to 50 percent. Koh can also look at acquisition and retention costs and leverage audience segmentation initiatives.

Finally, its new data architecture justified hiring a data scientist, who has built a machine learning model to predict customer demand.



OPTIMIZE RETENTION | VIEW THE FULL FUNNEL | PERSONALIZE THE CUSTOMER JOURNEY

Ritual

DATA STACK



Ritual

Ritual is a direct-to-consumer subscription wellness brand that allows customers to pause or cancel at any time — making customer retention a pillar of the company's long-term growth strategy. But its existing data architecture made it difficult to dig into the retention data and make informed business decisions.

Ritual hosted its site on Shopify, with a third-party plugin managing the subscriptions. Its ETL data pipeline was prone to failure, and the performance of its on-premise warehouse was steadily deteriorating. It relied on a nightly snapshot of the subscription, retention and

customer data stored in a warehouse, but couldn't analyze the numbers further.

So Ritual modernized its architecture, using Fivetran to centralize transactional data, email events, website interactions, ad data and post-purchase feedback into Snowflake. The Ritual team didn't need to build, configure or maintain the pipelines, which supplied analysis-ready SaaS data in near real time.

Here's how the new data stack, which includes Looker for business intelligence, has enabled Ritual to solve the retention puzzle:

- **Ad hoc data exploration.** Customer experts can explore the data themselves to answer questions like:
 - Does payment type impact retention?
 - Does exposure to an AB test pre-purchase impact retention?
 - Does a particular paid channel yield better-quality subscribers?
- **Personalizing the customer journey.** The data and lifecycle teams improved retention by identifying at-risk and high-value cohorts and testing different post-purchase experiences.
- **Visibility into the full funnel.** Users can see from the first website touch all the way through the subscription journey.

"The data showed that some customers were pausing very early in their subscription, so we sent them a personalized content series with tips and tricks on how to build a healthy habit, resulting in better retention and higher LTV."

– Brett Trani, Director of Data and Analytics, Ritual



Fivetran is the global leader in modern data integration. Our mission is to make access to data as simple and reliable as electricity. Built for the cloud, Fivetran enables data teams to effortlessly centralize and transform data from hundreds of SaaS and on-prem data sources into high-performance cloud destinations. Fast-moving startups to the world's largest companies use Fivetran to accelerate modern analytics and operational efficiency, fueling data-driven business growth.

For more info, visit [Fivetran.com](https://fivetran.com).