

Wyze increases conversions powered by complete customer profiles using RudderStack and Snowflake



Industry



Consumer electronics
Subscription services

Company size



101-250 employees

Headquarters



Kirkland, WA, USA

Highlights

- ▶ Wyze eliminated data silos by using RudderStack to build a warehouse native workflow centered around Snowflake
- ▶ Wyze leveraged RudderStack Profiles to automatically generate an identity graph, select and compute user features, and build a full view of the customer journey
- ▶ RudderStack Profiles eliminated a critical blind spot, allowing Wyze to track users as they move between the company's website, Amazon storefront, and different apps

Key stats

10x

increased productivity on
data engineering team

50+

user features automated
into a data pipeline in
hours, not weeks

3x

faster ML/AI project
deployments

3x

increase in marketing
campaigns driven by ML



RudderStack's warehouse native approach eliminated the manual processes bogging down our data engineers. With clean data at their disposal and automated workflows to route it downstream, they started providing our AI/ML, marketing, and product teams with actionable information to drive new models and power new processes."

Wei Zhou, Director of Data Engineering at Wyze

Overview

Founded in 2017 by three former Amazon employees, Wyze is on a mission to bring smart home technology to the masses. The company's first product, the Wyze Cam, launched at \$19.99, a "friend's price" ten times lower than most competitors' smart home security cameras. Selling directly to consumers online and relying on word-of-mouth and social media testimonials, the business grew from three to over 200 employees and was named one of Time Magazine's most influential companies of 2022.

With tens of millions of customers and a significantly higher number of IoT devices, Wyze has positioned itself as a leader in the field. To increase online sales, boost subscriptions, and fuel growth, Wyze's data engineering team tried to feed data to the Braze recommendation and personalization engine, hoping to steer customers toward relevant product purchases. However, the team encountered two significant hurdles during the process. There were too many event logs from disparate sources, and the data wasn't clean. These bottlenecks hindered progress for the team to deliver useful user features that leveraged event data, which impacted the rest of the business.

Without clean event data captured from the entire customer journey, the AI/ML teams relied on a smaller subset of data sources, leaving valuable data and insight out of the picture. Further downstream, the growth and customer acquisition teams had to deal with limited insights due to valuable customer journey touchpoints being left out of their analysis.

To further complicate the situation, Wyze was using Segment as its CDP, which was expensive, lacked robust consumer profile features, and created another data silo. This hampered their ability to drive their desired business outcomes - driving growth and customer LTV.

"Events were flooding in, but we couldn't extract any meaningful data," explains Wyze Director of Data Engineering Wayne Zhou. "We were losing the nuance and the detail, and everything was stuck in Segment's cloud instead of residing in our warehouse. We were spending a lot of money, and our data was siloed in an inflexible platform."

To remedy the situation, Wyze adopted RudderStack as its warehouse native CDP, using Event Stream, Reverse ETL, and JavaScript/Python Transformations features to automate ingesting, cleaning, and standardizing data before routing it from Wyze's Snowflake warehouse to downstream destinations.

It was a promising start, but RudderStack had even more in the pipeline - a tool that would fully resolve customer identities, and generate accurate 360 customer views. This exciting new product was RudderStack Profiles, which was in closed beta at the time, and Wyze was invited to join.

Challenge | Resolving identities to meet customers where they live

Wyze has two business areas, hardware sales and subscription services, which makes for a complex customer journey. The data engineering team needed to track users across multiple touchpoints, however, they lacked a 360 customer view.

"We couldn't see the entire customer journey," explains Zhou. "We could track users on our Shopify site and analyze their purchase and browsing histories, but they were invisible once they started ordering products and

services on our apps, after which point they rarely returned to the website. This blind spot prevented us from consolidating event data into cohesive user identities.”

Wyze lacked standardized data schemas and a way to join unique identifiers from these disparate sources. To simply add and map a new data source or user event required three data engineers and several days or weeks to complete.

These delays and deficiencies hobbled every team in the company, preventing data engineers, AI/ML specialists, and the marketing and subscription teams from doing their work. More importantly, they brought the development of the company’s Braze-powered recommendation engine to a standstill.

“We spent all our time matching customer behaviors to their identities,” adds Zhou. “Our engineers had access to different event logs but couldn’t tie everything together or discern the nuances in user actions to trigger personalized recommendations on all our platforms.”

Solution | Build complete profiles in Snowflake

Wyze used RudderStack Event Stream and Reverse ETL to standardize data collection and event schemas, eliminating their biggest obstacle to shipping a recommendation engine. Data engineers could finally consolidate clean, reliable data from different platforms and establish their Snowflake warehouse as a single source of truth.

This paved the way for Wyze to test RudderStack Profiles, a new product that tackles the complexities of identity resolution.

RudderStack Profiles makes it easy to build complete profiles in your warehouse. It automates the creation of an identity graph by extracting relevant data points across multiple sources, and generates the needed SQL code to create a unified table in Snowflake. Instead of spending hours and days authoring queries, Wyze’s engineers can easily select unique identifiers across hundreds of tables and run them through Profiles for accurate and complete identity graphs. They can then collect relevant user features to create 360 customer views to feed the company’s predictive analytics models.

“The Profiles pilot was an incredible success,” explains Zhou, “We have clean, consistent data across the user journey, including standardized schemas and identifiers. Profiles’ simplified workflows allow our data engineers to fully resolve identities, create user features, and present 360 customer views that are instantly operational in our ML models. It unlocks the clean data in our Snowflake warehouse delivered by RudderStack Event Stream, and frees our data engineers to route it downstream instead of prepping and sorting it.”

In combination with RudderStack Event Stream and Reverse ETL, Profiles has given Wyze ownership of its data by eliminating tedious and repetitive data cleaning, reducing irregularities from changing data sources, and speeding up onboarding new event streams. Wyze used these three tools to run an end-to-end customer data platform on top of Snowflake, accelerating the delivery of new features, and expediting the development of new activations.



When you have the power of RudderStack in hand, you can blast off right away. It’s so much easier to build a machine learning model once your designs are driven by clean data, useful user features, and 360 customer views.”

Wei Zhou, Director of Data Engineering at Wyze

Results | Unlock revenue growth with unified data




Wyze used RudderStack to build a streamlined warehouse native workflow that turbocharges activation. RudderStack Event Stream ingests clean user data from six sources into the company’s Snowflake warehouse. Then RudderStack Profiles takes over, resolving user identities and enabling the creation of user features and customer 360 views, which feed predictive models that compute actionable metrics like churn scores and customer lifetime values (LTVs). These metrics are then routed to the Braze personalization engine to orchestrate contextually relevant marketing and subscription campaigns.

Wyze is using customer 360 views to calculate propensity scores, tailor customer offers, and deliver personalized recommendations. The company’s AI/ML team is building actionable models to drive valuable insights. The company’s marketing team is leveraging this information to create campaigns, build canvases, and activate their customer data to drive revenue growth.

“RudderStack has given us the scale, scope, and speed to double and triple conversions because we’re pushing out 3x more scored campaigns.”

Zhou also estimates a 10x increase in productivity, especially in leveraging data in Wyze’s data pipelines to deliver new user features.

“RudderStack data can move in many directions, concludes Zhou. “We can send data to Braze, route it to our member services platform, or use it to trigger emails, text messages, and push notifications. We’re working on all kinds of stuff and boosting our data engineering and AI/ML teams’ horsepower to resolve identities faster and put identity graphs, user features, and 360 customer views into production faster.”

<div>Sources</div> <div>.Net; Android; Braze; iOS; Javascript; Python; Shopify; Snowflake; Webhooks</div>	<div>Destinations</div> <div>Snowflake, Braze</div>	<div>Warehouses</div> <div>Snowflake</div>
--	---	---