




WHITEPAPER

The Complete Guide to the Composable CDP



Contents

04	A Quick History of CDPs
06	Problems with Traditional CDPs
08	The Rise of the Composable CDP
10	Benefits of a Composable CDP
11	Composable CDP Use Cases
12	Getting Started with a Composable CDP



Preface

Organizations have more data than ever before at their disposal. Still, even with all of this information, many marketing teams find it difficult to drive outcomes that actually move the needle forward, and many data teams struggle to tie their work to revenue-facing functions.

For years, data and marketing teams alike have relied on traditional Customer Data Platforms (CDPs) to power personalization, targeting, and other customer-facing marketing efforts. However, with the massive adoption of cloud data warehouses, now any organization, regardless of industry or size, has access to near-infinite amounts of customer data. Many businesses are now adopting a Composable CDP to combat this challenge to activate the rich insights living in their data warehouse.

	“ We have to do things great for millions of customers every single day. We have to have technology to power that. That’s what we’re getting from the Hightouch ecosystem.  Bradley Breuer VP of Marketing
	“ The combination of Snowflake and Hightouch forms the cornerstone of our customer intelligence platform.  Tom Dinneny Customer Intelligence Platform Director
	“ One of the big value props of Hightouch is time-to-value. We’ve seen a lot of value very quickly.  Natalie Miles Head of Marketing Technology and Personalization

A Quick History of CDPs

A CDP is an all-in-one marketing and data platform designed to help companies collect and store customer data in a centralized database, so companies can easily build audiences, perform identity resolution, and sync that data to downstream operational tools.

All traditional CDPs operate as black boxes, offering little to no configurability when it comes to product features because everything is tightly integrated and bundled into a single product offering. As such, traditional CDPs have five core components:

Packaged CDP

 Event Collection	CDPs come equipped with proprietary SDKs (software development kits that organizations can use to collect behavioral events.
 Data Storage	CDPs offer fully managed storage to house customer data.
 Identity Resolution	CDPs have built-in identity resolution features to stitch together user actions and attributes across touchpoints.
 Audience Building	CDPs provide a rich suite of audience management tools to help marketers build user cohorts and orchestrate user journeys across marketing channels.
 Data Syncing	CDPs provide out-of-the-box integrations with hundreds of third-party APIs, enabling organizations to send audiences directly to their operational tools.


CDPs have been around for a long time, with Segment pioneering the space in 2011. Since then, data and marketing teams in almost every industry have relied on these platforms to power their customer-facing use cases.

These solutions emerged during a time when data collection, management, and activation were extremely challenging. Prior to CDPs, there was a huge rift between data and marketing teams because data tools weren't built for marketers, and marketing tools weren't built for data teams.

Moving data between various applications and systems was an engineering nightmare as it required a combination of custom scripts, API integrations, or manual CSV files. More often than not, the end result of these efforts was a spiderweb of pipelines powered by intricate workflows within various tools.

Traditional CDPs resolved this complex web of pipelines by introducing a single platform that automatically integrated with downstream applications. This meant data teams no longer had to worry about managing data flows, and marketing teams no longer had to worry about data accessibility. Data and engineering teams could implement event collection, and marketing teams could easily self-serve and build audiences to power their downstream digital channels.

Traditional CDP Architecture



Problems with Traditional CDPs

While CDPs set out to solve a key challenge, the tightly coupled and siloed architectures of these platforms introduce several complexities for today's modern organization.

Storage and Data Ownership

Traditional CDPs create a second source of truth because they force companies to store and manage data outside of their existing data infrastructure.

Flexibility

Legacy CDPs lock customers into an opinionated rigid data model. This makes it impossible to represent relationships between users and other entities like accounts, teams, workspaces, playlists, or devices.

Time-to-Value

The average CDP implementation takes over six months to complete, and that's not even accounting for onboarding time. Additionally, trying to support new use cases or adding new data sources introduces even more complexity because data teams are forced to build new data pipelines and rearchitect the CDP any time a new attribute needs to be stored.

Incomplete Data

Since CDPs only collect and store behavioral data, they create a fragmented view of the customer. Marketers have no ability to access or leverage any existing first-party attributes (income, age, gender, address) or the custom data models the data team has built. Some CDPs have tacked on features to interface directly with data warehouses, but these features are often lacking and separate from the core product offerings.


Cost

Traditional CDPs are extremely expensive because every feature is bundled together. This means companies have to pay for collection, storage, and modeling, even if they've already solved for these functions internally. Inevitably, many companies end up paying to store the same data twice (once in a warehouse and once in a CDP.)

Ultimately traditional CDPs come in a “one-size-fits-all” model that makes them very rigid for hyper-specific use cases. For example:

- With a traditional CDP, it’s **possible** to send an email to cart abandoners, but it’s almost **impossible** to send that same email with a targeted local store incentive because store data is not natively available in a CDP.
- With a traditional CDP, it’s **possible** to show ads to people who have visited product pages, but it’s almost **impossible** to ensure those same ads aren’t shown when a product is out of stock. Keeping product catalogs up-to-date in ad platforms is not possible with a CDP because data has to be conformed under a strict user/account object, and product data is not natively available.
- With a traditional CDP, it’s **possible** to send an SMS about a new promotion, but it’s almost **impossible** to send that same SMS to people who have a high propensity of redeeming it. A CDP cannot power recommendation engines or data science models. The data warehouse is the only platform that has the data needed to facilitate these use cases.

Each of these marketing use cases points out a fundamental problem: traditional CDPs have no understanding of custom objects (e.g., local stores, products, propensity models). With a traditional CDP, it’s impossible to create a Spotify Wrapped campaign or serve personalized recommendations to specific audiences. Traditional CDPs boil everything down to clicks and sessions.



The reality is businesses are much more complicated than behavioral data, and modern marketers should have access to a 360-degree view of the customer. That includes everything from behavioral data to offline actions, to data science models, to first-party attributes. This level of granularity is only possible with a Composable CDP powered by the data warehouse.

The Rise of the Composable CDP


A Composable CDP has the same goal as a traditional CDP: **activating customer data to your downstream tools**. But rather than operating as a separate entity and storing data outside of your current data infrastructure like a traditional CDP does, a Composable CDP lets you create audiences, personalize journeys, and send your existing customer data directly from your data warehouse to your frontline marketing tools.

While CDPs set out to solve the correct problem, they did so with the fundamental flaw that they did not build upon an immensely powerful technology rapidly being adopted by organizations as their central store of customer data – the cloud data warehouse.

The rapid adoption and investment in cloud data warehouses has caused many companies to realize that they don't need to purchase a separate platform to store data they already own in their warehouse—inevitably leading many companies to build on top of their existing data infrastructure and adopt a Composable CDP.


By simply integrating Hightouch with their existing data warehouse, organizations are able to form what has since been coined the “Composable CDP” by the likes of cloud leaders like [Snowflake](#) and [Databricks](#). This architectural approach provides unparalleled flexibility allowing companies to solve complex use cases in both a timely and cost-efficient manner, enabling organizations to own and act on all of their data.

Composable CDP Architecture



One of the biggest differences between a traditional CDP and a Composable CDP is that a Composable CDP lets you start with your current data (wherever it is) and activate it immediately rather than implementing and purchasing an entirely new tool.

Traditional CDP vs. Composable CDP



The modularity of this architecture gives companies far more extensive capabilities and flexibility to uniquely solve their most complex customer-facing use cases. This means organizations don't have to conform data to the requirements and constraints of another platform and can instead take advantage of their existing data assets to drive value immediately.

The easiest way to understand this difference is to view them as two restaurant menu offerings.

Traditional CDPs follow a table d'hôte model, which refers to a fixed menu option where the entire meal is offered at a set price. Conversely, a Composable CDP follows an à la carte model, with individual pricing where people can choose items tailored to their specific preferences.




We didn't choose an off-the-shelf CDP like Segment because a substantial amount of engineering is required to get it up and running, and at the end of the day, it is just a second source of truth. Why would we buy a CDP when all of our data modeling is already being done directly in Snowflake?



Ryan Newsome
Head of Data & Analytics

[Read full story >](#)

Benefits of a Composable CDP



Flexibility

A Composable CDP gives organizations immediate access to any and all of their data—not just clickstream events. Ultimately this means marketing teams can leverage any data type to easily accommodate the complex breadth of their use cases.



Security

Because a Composable CDP simply builds on top of existing infrastructure, organizations don't have to worry about data leaving their cloud environment.



Modularity

Companies have full control to choose what technologies and processes they use for data collection, storage, modeling, and activation. This allows organizations to tailor their architecture to the specific organizational outcomes they're looking to drive.



Time-to-Value

A Composable CDP enables organizations to circumvent long implementation times and activate their existing data assets immediately to drive business outcomes.

Composable CDP Use Cases

This modular architecture works with any organization regardless of maturity, and businesses in all industries are seeing the immediate benefits of building on top of their data warehouse.


  IntelyCare Saves \$1M in Marketing Spend	“ We ran an experiment around first-shift bonuses and followed different test groups through the application process to see how much the bonus influenced shift completion rate. Without Hightouch, this experiment cost us \$1.3M. With Hightouch we're able to intelligently target applicants and achieve the same incremental gains at an annual cost of \$260k.  Ben Tenglesen VP of Data Science Read full story >
  Chalhoub Group Increases Ad Revenue by 40% and Lowers CAC by 30%	“ We used to have to go to our CRM and spend hours building customer segments, and manually upload them to our marketing platforms, but this data wasn't refreshed, clean, or debugged in any way. We could only do this once a year because we couldn't afford to take two weeks off to upload static files.  Tomasz Mazur Head of Customer Acquisition Read full story >
  Warner Music Group Powers 1000+ custom audiences	“ Within six weeks, we implemented Hightouch on Snowflake and got out hundreds of audiences, which made marketers and agencies happy.  Itay Rahat VP of Product, Audiences, and Marketing Technology Read full story >

Getting Started with a Composable CDP

Most organizations think they need to reach a certain level of data maturity before they can fully take advantage of a Composable CDP. In reality, this is a misconception. Hightouch offers the capabilities of a traditional CDP without the tightly coupled architecture.


Customer 360 Toolkit

Stitch your existing customer data into rich, actionable profiles directly in your warehouse. Link all of your user actions and attributes to a single comprehensive profile without writing any of code.




Activation

Automate and schedule data syncs to 200+ destinations to ensure your data is always fresh and up-to-date, so your marketers can run ad-hoc experiments and campaigns in minutes.




Customer Studio

Build and define rich audience cohorts to power your most complex personalization use cases. Perform A/B testing and cross-channel orchestration and easily analyze audience breakdowns and performance.



Your favorite companies use Hightouch to sync data securely at scale.



[Schedule a demo](#)

[Create a free workspace](#)