ANALYST CONNECTION

Sponsored by: Twilio Inc.

Organizations must consider a number of factors when selecting a customer data platform, exploring both hybrid and warehouse-native approaches for business success.

Unlocking CX Success: Making the Right CDP Architecture Choice

September 2025

Questions posed by: Twilio

Answers by: Tapan Patel, Research Director, AI-Enabled Customer Data and Analytics

Q. How do use case requirements and prebuilt functionality need to guide customer data platform (CDP) architecture decisions?

A. A successful CDP implementation rests on selecting an architecture that meets an organization's immediate needs and scales with its evolving marketing and customer experience (CX) use cases, skills, and martech ecosystem. A production-ready data foundation is key to a successful cross-functional customer data initiative, as it minimizes risks by ensuring data quality, consistency, and accessibility. Organizations must assess how a CDP architecture contributes to or detracts from these foundational strengths. They face a choice between warehouse-native CDP or hybrid CDP approaches. Each has distinct advantages and disadvantages depending on the specific context.

Hybrid CDPs provide prebuilt functionality that lessens reliance on IT and data engineering resources, helping reduce the burden for certain predefined tasks like data ingestion, identity resolution, profile management, and data activations. These solutions offer integrated capabilities for common CDP tasks, with no-code user interfaces (UIs), connectors, and data models for batch or real-time use case activations. However, hybrid CDPs may create dependencies, and organizations may require ongoing coordination between prebuilt CDP capabilities and integration with warehouse infrastructure.

Warehouse-native CDPs offer flexibility by leveraging a company's existing cloud data infrastructure as the foundation while allowing organizations to select specialized components for specific CDP functions. This approach may work well for companies with strong data and IT capabilities and existing cloud data warehouse investments, as it can enable quick deployment of initial use and provide customization opportunities. However, internal expertise is needed to manage data pipelines, profile creation, and manage these integrations.

The critical factor is ensuring the chosen architecture supports multifunctional use case activation across marketing, customer service, sales, commerce, and analytics teams. The choice depends on martech maturity, data and IT skills and resources, and the complexity of cross-departmental requirements. Each approach delivers distinct forms of value.

CDPs with hybrid architectures typically provide predictable multi-departmental activation through integrated customer dataflows and prebuilt capabilities. Warehouse-native approaches offer data control and customization potential aligned with organizations' existing data and IT resources and infrastructure investments.

Q. Beyond licensing, how can an organization assess the true cost and resource commitment of a CDP?

A. Initial licensing costs often hide the true expense of CDP ownerships, which can significantly impact long-term value. Organizations frequently underestimate ongoing expenses including implementations, specialized head count, operational resources, and the costs of expanding into new CDP use cases.

With warehouse-native CDPs, consider the shift of financial costs upstream to the organization's cloud data warehouse, which can lead to increased compute expenses. This reallocation requires a dedicated internal data and IT team to manage CDP components, from data aggregation to activation. Organizations must weigh the resources needed to build, maintain, and support data pipelines, integrations, interoperability, and user requests and assess whether this will slow down new CDP use case activation. These resource and operational decisions should account for the opportunity costs of data and IT teams, as what seems like a lower initial cost can become unpredictable long-term expenses if not carefully managed.

Hybrid CDPs provide a viable option for organizations prioritizing a predictable cost model, minimizing dependency on IT and data teams, and access to prebuilt features. While initial licensing costs can be higher, they typically encompass a wide array of prebuilt functionalities, such as data ingestion, identity resolution, segmentation, and activation. This translates to a lower initial operational burden, as the vendor manages the underlying infrastructure and maintenance. For businesses with less mature data teams or limited IT resources, the reduced complexity and deployment cycles of these hybrid CDPs can also help accelerate the activation of new use cases. While there might be less granular control for some hybrid CDPs or a need to adapt to the vendor's data model, the lower demand on internal resources can lead to more manageable long-term costs.

Ultimately, the optimal CDP cost model aligns with an organization's strategic priorities, technical capabilities, and resource availability. A thorough evaluation of the total cost of ownership — encompassing both up-front and ongoing operational demands — is crucial for sustainable customer data activation.

Q. How can CDP architectures be leveraged for real-time use cases, given their capabilities, limitations, and promise?

A. Real-time capabilities represent a technically demanding aspect of CDP implementations, necessitating the evaluation of streaming data ingestion through to activation at an enterprise scale. True real-time functionality extends beyond simple data ingestion and aggregation to encompass immediate identity resolution, profile unification, and subsecond activation across multiple channels.

Hybrid CDPs offer mature real-time capabilities and infrastructure to process high-volume data and customer interactions during activities like site personalization or new purchase decisions while maintaining data quality and consistency. Warehouse-native CDPs face inherent challenges with real-time processing due to the batch-oriented nature of traditional data warehouses. While these systems are evolving to support real-time interactions, achieving true real-time processing often requires assembly and infrastructure layers that introduce complexity and potential failure points.

Building enterprise-grade, real-time capabilities demands caching mechanisms, reliability protocols, and observability frameworks, which many organizations lack the expertise to build effectively. Data observability is critical for the success of any real-time data initiative, enabling teams to understand the state and health of their data, pipelines, and infrastructure. Hybrid CDPs should offer these as features, simplifying management and providing a more unified approach to system health and compliance. Warehouse-native solutions require selecting and integrating separate tooling for each aspect, placing a burden on the internal team for overall system health, performance, and compliance.

Simply having real-time data ingestion is insufficient; organizations must assess whether their teams possess expertise to build and maintain sophisticated real-time systems or whether partnering with CDP vendors would accelerate time to value while reducing technical risk. Companies should focus on understanding end-to-end latency requirements for critical use cases, testing vendor claims with realistic data volumes, and ensuring the architecture chosen supports the reliability standards required for customer-facing real-time applications.

Q. How important is it for CDPs to empower marketing to segment customers, orchestrate journeys, and activate data without relying on specialized data or IT teams?

A. When marketing and customer experience teams can independently create segments, launch campaigns, and activate journeys, organizations gain the ability to be proactive and react quickly to market opportunities and customer behavior changes.

No-code capabilities vary significantly across CDP solutions, ranging from basic drag-and-drop segmentation tools to advanced workflow builders enabling complex customer journey orchestration. Hybrid CDPs typically provide marketer-friendly interfaces with visual builders and prebuilt connectors. Warehouse-native CDP solutions, despite improvements in user interface design, still require data-savvy users for data pipelines and modeling. However, a successful implementation extends beyond UI and no-code capabilities. It also encompasses organizational readiness for data-driven decision-making, adapting marketing and CX workflows accordingly, and ensuring the selected CDP empowers its target users, rather than adding complexity.

A CDP that is a mismatch with the marketing or CX team's skills and capabilities can lead to shelfware, underutilization, IT dependency, and a failure to quickly realize value. Consider whether the organization needs basic segmentation capabilities or requires sophisticated features like real-time trigger management and cross-channel orchestration. Companies with limited technical resources or those prioritizing business user autonomy will find value in hybrid CDPs offering self-service functionality with appropriate guardrails, audit trails, and governance controls. Those with

established data engineering and IT teams that are comfortable with additional custom effort for real-time use cases may prefer the warehouse-native CDP route.

Q. What factors should guide organizations in selecting prebuilt AI models and workflows versus building custom ones as part of CDP and AI evaluations?

A. AI integration in CDP platforms presents organizations with a choice: leverage prebuilt AI models for speed or build custom data science workflows for deeper customization. Organizations must align their CDP selection with a realistic appraisal of data and AI skills, the need for AI functionality, and their operational maturity level. Warehouse-native CDPs enable custom predictive AI model development through direct cloud data warehouse access to avoid complex data movement and gain flexibility. Hybrid CDPs offer prebuilt predictive AI models that can be customizable and integrated into cloud data warehouses to reduce implementation time.

Organizations that have in-house data engineering and data science skills and have invested up front in data infrastructure will find value in a warehouse-native CDP for AI use cases. Hybrid CDPs can handle trade-offs between speed and customization by offering prebuilt AI models for timely deployment alongside the ability to update these models based on data in cloud warehouses for real-time use cases. The key is to honestly assess the organization's current data and AI maturity and use case needs and where it wants to be in the future.

Early stage AI adopters may gain more value from prebuilt models that quickly activate data for marketing, while organizations with data science teams can prioritize custom model development. In both cases, success depends on high-quality, well-governed data and the readiness to adopt AI-driven insights. Without accurate and contextualized data, even advanced algorithms produce unreliable outcomes. The most effective implementations match warehouse-native or hybrid CDP's AI capabilities to organizational maturity while allowing flexibility as teams grow.

About the Analyst

Tapan Patel, *Research Director, AI-Enabled Customer Data and Analytics*

Tapan Patel is research director for the Customer Data Platform and Artificial Intelligence and Analytics Software market segments and a member of the Customer Experience research team at IDC. Tapan's core research coverage includes market trends, end-user requirements, use cases, market sizing, and business models for these critical segments. He is lead analyst for the CDP and Customer Analytics markets, used by brands to improve customer engagements, journeys, and personalization across all touch points.

MESSAGE FROM THE SPONSOR

Twilio Segment is a hybrid CDP that offers real-time event computation and activation alongside robust warehouse integration, allowing businesses to innovate quickly with a composable architecture that adapts as customer needs change.

Supporting real-time engagement for marketing and product teams, and centralized control and enrichment for data teams, Twilio Segment empowers businesses to maximize data warehouse investments and drive dynamic, personalized customer engagement.

Twilio Segment CDP offers:

- Real-time event and warehouse activation
- Built-in identity resolution and data governance
- Zero-copy architecture
- No-code audience building for marketers
- Extensibility for data teams

Learn more: https://segment.com/composable-cdp/.

IDC Custom Solutions

IDC Research, Inc.
140 Kendrick Street
Building B
Needham, MA 02494
T 508.872.8200
F 508.935.4015
blogs.idc.com
www.idc.com