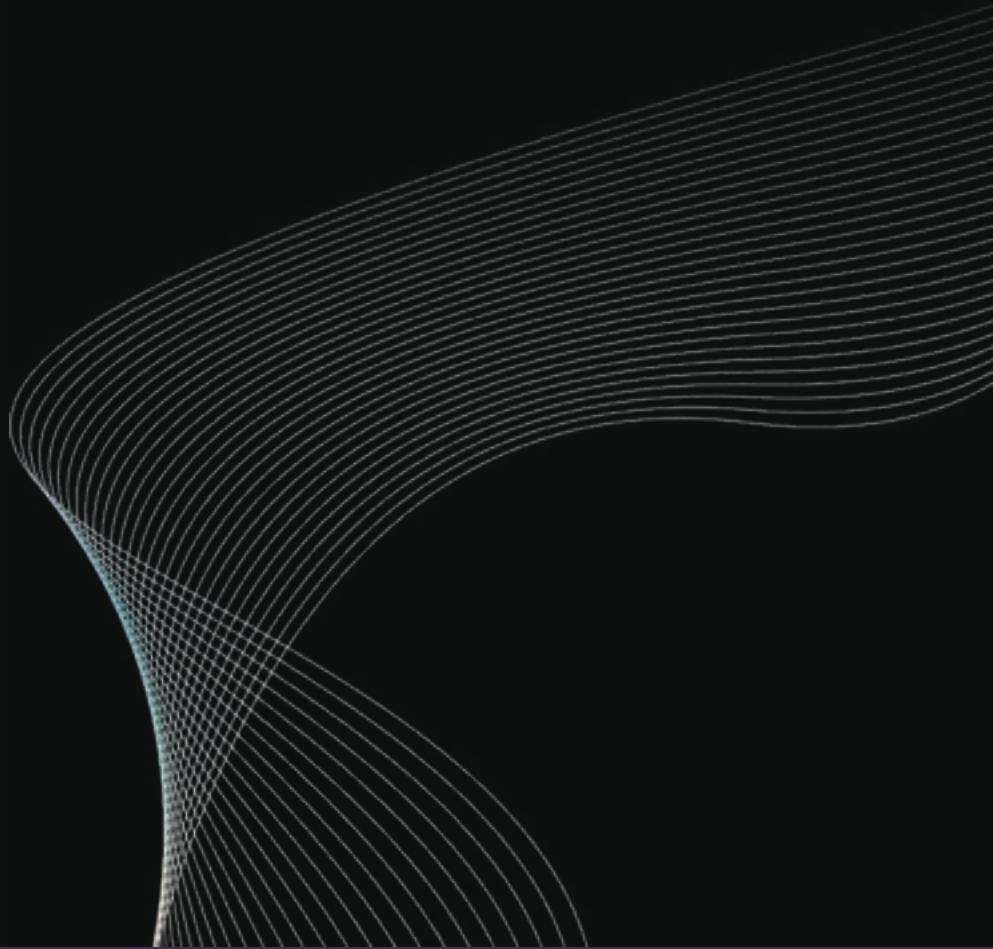




October 2025 MarTech CDPs

VENDOR PULSE REPORT



Contents

Executive Snapshot

3

Retail & eCommerce Priorities

5

Market Alignment with Leader Priorities

8

Vendor Maps: Archetypes and Capability Lenses

14

Vendor Deep Dives

26-160

About VOX

161



Executive Summary

Strategic Insights for Enterprise Retail & eCommerce Leaders

Welcome to the VOX Pulse Report, your guide to the latest strategic priorities and vendor alignment insights directly from enterprise retail and eCommerce leaders navigating today's Customer Data Platform (CDP) landscape. Drawing directly from recent VOX data, **this summary outlines the targeted outcomes senior executives prioritize, the critical capabilities they urgently require, and how effectively current vendor solutions meet these dynamic operational needs.** Use this summary to quickly orient yourself within the market, identify key strengths or gaps among vendors, and navigate deeper into detailed vendor evaluations.

Key Outcomes Leaders Aim to Achieve

Recent VOX data reveals senior executives in retail and eCommerce are highly focused on delivering personalized, real-time customer experiences that boost conversions and drive loyalty without compromising customer trust or privacy. With diminishing third-party data availability and rising customer acquisition costs, leaders increasingly prioritize first-party data to drive precise merchandising and impactful customer interactions. Today's retail leaders explicitly define real-time activation as instantaneous responses within active customer sessions, underscoring a clear shift from delayed insights to immediate, actionable engagement.

Operational pressures heighten the urgency around these priorities, prompting investments in reliable, real-time data infrastructure, seamless multi-channel integration, unified customer profiles, and embedded AI-driven analytics. Compliance, privacy, and data security have also become integral components of broader customer experience strategies, reflecting a firm organizational commitment to governance excellence.

For an in-depth exploration of these leader priorities and their strategic context, see [Retail & eCommerce Priorities: What's Urgent for Leaders Now \(page 6\)](#).

Critical Capabilities Leaders Demand

To translate these strategic outcomes into execution, leaders emphasize a defined set of essential capabilities. These include a **reliable, real-time data infrastructure** featuring **continuous observability, proactive anomaly detection**, and flexible analytics tools such as **direct SQL or lake-house access**. **Seamless integration across commerce, marketing, and service platforms** through pre-built connectors is now table-stakes for efficient, bi-directional data flow.

Leaders further highlight the importance of a **unified customer identity across online and offline interactions** for consistent customer experiences. Additionally, **real-time activation capabilities**—particularly **sub-second decision-making for instant engagement**—and **embedded predictive analytics** are prioritized to drive targeted merchandising and personalized recommendations. Robust built-in **governance, privacy compliance, and consent management capabilities** are now foundational expectations.



Market Landscape at a Glance

VOX data illustrates significant maturity and alignment among CDP vendors with retail and eCommerce leader priorities. Nearly all evaluated vendors (31) address at least 10 of the 16 critical capabilities identified by leaders. Vendors such as [BlueConic](#), [Salesforce Data 360](#), [Algemony](#), [RudderStack](#), and [Treasure Data](#) exhibit broad alignment across multiple priority areas. Foundational capabilities including deterministic identity resolution, native inbound connectors, predictive analytics, unified event collection, sub-second decision-making, fast inventory synchronization, and direct SQL access are widely supported. However, **gaps remain in capabilities such as real-time observability, comprehensive data lineage, and proprietary analytical model deployment.**

Real-time/Edge emerges as the most cross-cutting strength, consistently appearing across all five vendor archetypes, though the underlying reasons vary significantly. For Engagement Hubs and Commerce Engines, real-time capabilities emphasize activation speed within customer interactions; Customer Data Infrastructure vendors highlight low-latency data pipelines; and Suite CDPs provide policy-safe real-time capabilities within compliance-centric frameworks. Leaders commonly pair specific vendor archetypes with corresponding capability lenses (Data Visibility & Quality, Security, Consent & Residency, Customization & Composability, In-Session Activation) to validate alignment with their distinct operational priorities

Insights by Vendor Archetype

Distinctive patterns emerge across vendor archetypes aligning with different strategic requirements:

- **Independent CDPs** such as BlueConic, Uniphore, and Amperity provide specialized data unification, segmentation, and real-time integration capabilities, with strong technical openness. [\[Explore the full Independent CDP coverage map on page 17\]](#)
- **Composable CDI platforms**, notably RudderStack and Treasure Data, offer significant flexibility with warehouse-first strategies and robust real-time integrations, aligning well with composable infrastructure priorities. [\[Explore the full Composable CDI coverage map on page 19\]](#)
- **Engagement Hubs**, prominently represented by Oracle Unity CDP and Zeta Global, excel in journey orchestration and customer engagement, focusing heavily on real-time decision-making and activation. [\[Explore the full Engagement Hub coverage map on page 20\]](#)
- **Commerce and Personalization Engines** including Algemony, Bloomreach, and SAP Emarsys deliver robust real-time personalization and merchandising solutions, ideally suited for immediate customer interactions. [\[Explore the full Commerce and Personalization Engine coverage map on page 18\]](#)
- **Integrated Suite CDPs** such as Salesforce Data Cloud and Adobe Experience Platform stand out for their comprehensive marketing integration, governance, and compliance capabilities. [\[Explore the full Suite CDP coverage map on page 21\]](#)



Insights by Capability Lenses

Vendor performance across key capability lenses prioritized by leaders reveals distinct strengths and notable gaps:

- **Data Visibility & Quality:** Consistent vendor strengths are evident in live data lineage, direct SQL access, and automated anomaly quarantine, with only select vendors not fully meeting specific requirements. [[Explore the complete Data Visibility & Quality capability map on page 22](#)]
- **Security, Consent & Residency:** Vendors show substantial alignment in role-based access controls (RBAC), instant consent suppression, and regional data residency compliance, reflecting high industry maturity in these critical governance areas. [[Explore the full Security, Consent & Residency capability map on page 23](#)]
- **Customization & Composability:** Notable capability gaps exist around uploading proprietary analytical models and providing direct SQL access, indicating areas leaders must carefully assess. [[Explore the complete Customization & Composability capability map on page 24](#)]
- **In-Session Activation:** About half of the vendors demonstrate robust capabilities for real-time, personalized activation such as Unified Event Collection, Fast Inventory Sync, Sub-Second Decision APIs, Audience Syndication, and Real-Time Profile APIs, highlighting significant alignment as well as some key gaps. [[Explore the full In-Session Activation capability map on page 25](#)]

For a detailed exploration of overall vendor capabilities and their alignment with retail and eCommerce leaders' strategic priorities, refer to the [Market Alignment with Leader Priorities](#) chapter (page 8).

Vendor Deep Dives by Archetype

To review specific vendors and explore in-depth how each aligns with the priorities and preferences of retail and eCommerce leaders, select any vendor below:

- **Independent CDP Vendors:** [Amperity](#), [BlueConic](#), [Contentstack](#), [Redpoint](#), [Session AI](#), [Uniphore](#) ([starts at page 27](#))
- **Composable CDI Vendors:** [Netcore](#), [RudderStack](#), [Tealium](#), [Treasure Data](#), [Twilio Segment](#) ([starts at page 73](#))
- **Engagement Hub Vendors:** [Blueshift](#), [Braze](#), [Cordial](#), [Dotdigital](#), [FirstHive](#), [Listrak](#), [Marigold](#), [Microsoft Dynamics 365 Customer Insights CDP](#), [Ometria](#), [Optimove](#), [Oracle Unity CDP](#), [Zeta Global](#) ([starts at page 94](#))
- **Commerce and Personalization Engine Vendors:** [Algonomy](#), [Bloomreach](#), [Iterable](#), [SAP Emarsys](#) ([starts at page 54](#))
- **Suite CDP Vendors:** [Adobe Experience Platform](#), [mParticle](#), [ReactorData](#), [Salesforce Data Cloud](#) ([starts at page 144](#))

Retail & eCommerce Priorities: What's Urgent for Leaders Now

This chapter objectively captures the current priorities retail and eCommerce leaders explicitly identify as most urgent and strategically critical based on VOX's latest preference data. It highlights key areas leaders prioritize for immediate operational attention and explores how emerging market pressures shape their strategic decisions. The insights presented here strictly reflect leader-defined preferences, challenges, and recommended focus areas, providing clarity around what matters most today in retail and eCommerce.

Outcomes Leaders Target

Leaders are explicitly prioritizing personalization that boosts conversions and loyalty, without compromising customer trust or privacy. Real-time engagement is now essential, shifting focus from delayed reporting to instant, actionable insights. As third-party data fades and acquisition costs rise, first-party data becomes central, fueling targeted merchandising and customer interactions.

Critical Capabilities Leaders Need Today

Reliable, Real-Time Data Infrastructure

Leaders require comprehensive visibility into data streams with real-time lineage tracking and automated anomaly detection. Direct SQL or lake-house analytics capabilities are preferred for flexibility over restrictive platforms.

Effortless Integration Across All Channels

Seamless, bi-directional integration across commerce, apps, customer service, and advertising channels through pre-built, easy-to-deploy connectors is now baseline.

Unified Identity Across Channels

Unified customer profiles that accurately bridge online identifiers with offline interactions are foundational for consistent and reliable customer experiences.

Genuine Real-Time Activation

Retailers explicitly define real-time as instant responses during active customer interactions, emphasizing sub-second decisions, rapid inventory updates, and immediate content personalization.

Advanced Analytics & Predictive Capabilities

Leaders increasingly rely on AI-driven analytics and built-in predictive capabilities, demanding structured, AI-ready data alongside seamless integration for proprietary analytical models.

Built-In Governance, Privacy & Compliance

Detailed role-based access controls, immutable audit trails, immediate consent management, and stringent regional data compliance have become critical expectations.



Key Capability Lenses Leaders Prioritize

When evaluating vendors, retail and eCommerce leaders consistently highlight four primary capability lenses as critical. These lenses guide leader evaluations and heavily influence their vendor selections.

- **Data Visibility & Quality:** Real-time monitoring of data streams with detailed lineage tracking and proactive quality controls, ensuring accuracy and trustworthiness.
- **Security, Consent & Residency:** Robust governance features, including field-level access controls, immediate consent enforcement, and regional data residency compliance.
- **Customization & Composability:** Flexibility to integrate proprietary models and existing analytics infrastructure seamlessly, minimizing dependency on external professional services.
- **In-Session Activation:** Real-time personalization and decision-making capabilities, essential for immediate customer interactions and operational responsiveness.

Why These Priorities Are Urgent

Leaders face acute operational pressures:

- **Multiple Data Sources:** Fragmented data slows operations, intensifying demand for unified, accurate data views.
- **Increasing Integration Complexity:** Growing costs and unreliability of custom integrations push teams toward native, pre-built solutions.
- **Rising Importance of First-Party Data:** Declining cookie viability makes efficient first-party data management mandatory.
- **Operational Real-Time Expectations:** Explicit requirements for instantaneous customer interactions redefine vendor selection criteria.
- **Essential Predictive Merchandising:** Predictive analytics directly influence merchandising decisions, shifting from optional to essential.

Board-Level Privacy and Ethics: Privacy compliance is now non-negotiable, necessitating demonstrable compliance with stringent standards.

Leader-Derived Strategic Recommendations

Insights from retail and eCommerce leaders suggest effective personalization fundamentally relies on robust and flexible data infrastructures. Leaders targeting immediate, in-session customer interactions prioritize vendors proven in rapid event handling, real-time inventory synchronization, and instantaneous personalization capabilities. Conversely, leaders focused heavily on analytics flexibility and data warehousing prioritize evaluating vendors that support open integration and deployment of proprietary predictive models.

Ultimately, the data indicates leaders consistently demand clear, demonstrable vendor capabilities, ensuring vendor promises align closely with actual operational performance. Retail and eCommerce leaders have decisively aligned their investments toward reliable, flexible, and privacy-compliant data solutions.

Market Alignment with Leader Priorities

By closely mapping enterprise retail and eCommerce leaders' latest priorities and preferences with comprehensive vendor evaluations, VOX highlights the evolving maturity among Customer Data Platform (CDP) vendors. The current landscape shows significant vendor capability alignment, with nearly all evaluated vendors addressing at least 10 of the 16 critical capabilities identified by leaders.

Several vendors, including BlueConic, Salesforce Data 360, Algonomy, RudderStack, and Treasure Data, closely align with these leader-driven priorities, showcasing notable coverage across multiple capability areas. Their alignment reflects current strategic objectives emphasized by retail and eCommerce leaders, highlighting broad industry alignment rather than individual winners.

Diverse Approaches in the Market

Based on the latest insights and preferences captured by **VOX** from enterprise retail and eCommerce leaders, executives have shown interest in exploring a wide variety of vendor archetypes—suggesting that multiple strategic approaches are viable depending on organizational objectives. These approaches broadly include:

Independent CDPs, emphasizing specialized data unification and segmentation capabilities;

»» **Vendors:** Amperity, BlueConic, Contentstack, Redpoint, Session AI, Uniphore.

Composable Customer Data Infrastructure (CDI), which provides flexibility around data warehouses

and real-time streaming;

»» **Vendors:** Netcore, RudderStack, Tealium, Treasure Data, Twilio Segment.

Engagement Hubs, highlighting robust journey orchestration and customer engagement capabilities;

Vendors: Blueshift, Braze, Cordial, Dotdigital, FirstHive, Listrak, Marigold, Microsoft Dynamics 365

»» Customer Insights CDP, Ometria, Optimove, Oracle Unity CDP, Zeta Global.

Commerce and Personalization Engines, strongly aligned with real-time, personalized customer interactions;

»» **Vendors:** Algonomy, Bloomreach, Iterable, SAP Emarsys.

Integrated Suite CDPs, offering comprehensive marketing and governance capabilities;

»» **Vendors:** Adobe Experience Platform, mParticle, ReactorData, Salesforce Data 360.

Collectively, these archetypes demonstrate two primary patterns: organizations prioritizing

composable, flexible infrastructure, and those **emphasizing integrated activation capabilities**.



Market Strengths and Opportunities

When mapping retail and eCommerce leaders' current priorities and requirements against publicly available vendor documentation and capability statements at scale, VOX has identified clear areas of vendor strength and specific gaps requiring careful consideration.

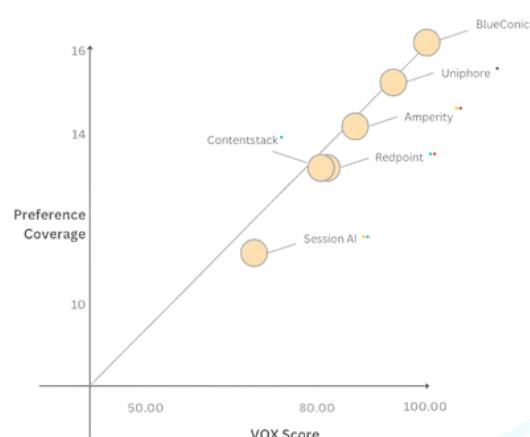
Retail and eCommerce leaders have strongly standardized around foundational features such as deterministic identity resolution, native inbound connectors, and embedded predictive analytics—all three capabilities are universally available (**100% vendor coverage**). Similarly, unified event collection (**88% coverage**), sub-second decision APIs (**85% coverage**), fast inventory synchronization (**82% coverage**), and direct SQL access (**82% coverage**) are well-established, indicating robust maturity in core data management and real-time operational capabilities.

However, retail and eCommerce leaders with advanced or specialized requirements should closely examine certain capability gaps. Real-time observability and comprehensive data lineage, essential for ensuring trustworthy data flow, have significantly lower availability (**52% vendor coverage**). Additionally, the flexibility to deploy proprietary analytical models remains limited (**36% vendor coverage**), representing a notable limitation for organizations prioritizing customized analytics.

Practically speaking, leaders should validate these lower-coverage capabilities against their strategic objectives and operational needs during vendor selection to avoid unexpected functional gaps.

Observations by Vendor Archetype

When evaluating vendor archetypes against the specific priorities of retail and eCommerce leaders, certain distinctive patterns and strengths emerge. Independent CDPs offer strong technical openness, composable CDIs excel in flexibility and real-time integration, Engagement Hubs focus prominently on journey orchestration, Commerce and Personalization Engines deliver robust real-time personalization capabilities, and Suite CDPs provide comprehensive governance and compliance functionalities.



Independent CDPs generally provide strong overall coverage, with vendors like BlueConic and ActionIQ (Uniphore) showcasing robust technical openness and real-time capabilities that align closely with leader-defined priorities.

[Explore the full Independent CDP coverage map on page 17]

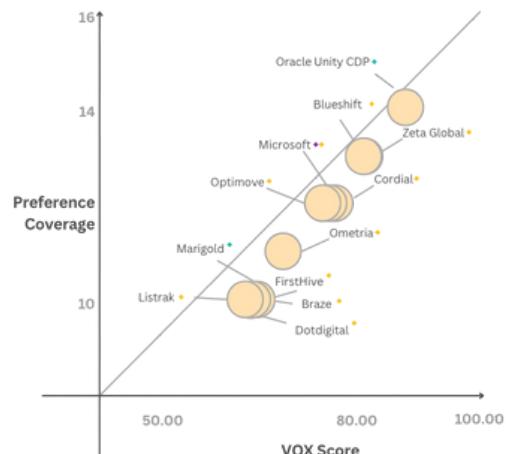


Observations by Vendor Archetype



Composable and CDI platforms, notably RudderStack and Treasure Data, consistently deliver flexibility in warehouse-first strategies and real-time integrations, aligning well with organizations prioritizing composable infrastructure.

[\[Explore the full Composable CDI coverage map on page 19\]](#)

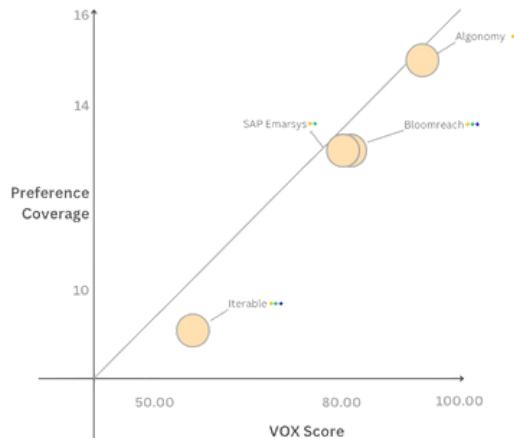


Engagement Hubs, with Oracle Unity CDP standing out, primarily emphasize journey orchestration capabilities, with varying degrees of real-time decision-making features.

[\[Explore the full Engagement Hub coverage map on page 20\]](#)

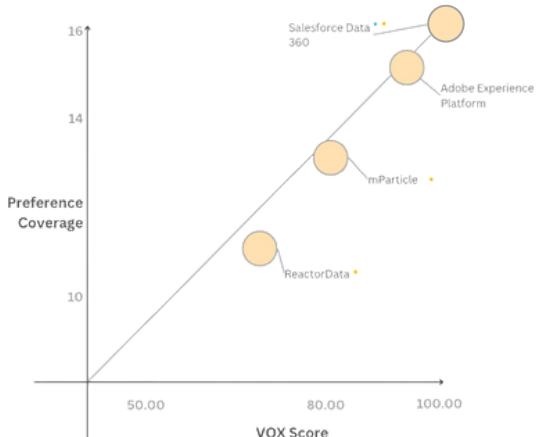


Observations by Vendor Archetype



Commerce and Personalization Engines, particularly Algonomy, Bloomreach, and SAP Emarsys, demonstrate strengths in real-time personalization and merchandising, catering to immediate customer interaction scenarios.

[\[Explore the full Commerce and Personalization Engine coverage map on page 18\]](#)



Suite CDPs such as Salesforce Data 360 and Adobe Experience Platform exhibit robust governance and compliance features suitable for enterprises prioritizing integrated and comprehensive compliance solutions.

[\[Explore the full Suite CDP coverage map on page 21\]](#)



Observing Capabilities Through Key Lenses

Having reviewed high-level market capability coverage, we now delve into specific capabilities that matter most to retail and eCommerce leaders. By mapping leader priorities against detailed vendor documentation and public information, VOX has surfaced distinct vendor strengths and critical gaps through several key capability lenses:

Data Visibility & Quality

Retail and eCommerce leaders require clear visibility and proactive quality management for their data. The heatmap indicates strong and consistent coverage across live data lineage, direct SQL access, and automated anomaly quarantine among 14 vendors, with only select vendors not meeting some key requirements.

[\[Explore the complete Data Visibility & Quality capability map on page 22\]](#)

Security, Consent & Residency

Governance capabilities, including role-based access controls (RBAC), instant consent suppression, and data residency compliance, are crucial for retail and eCommerce leaders. The heatmap shows similarly strong coverage across these areas among 16 vendors, indicating substantial alignment from vendors.

[\[Explore the full Security, Consent & Residency capability map on page 23\]](#)

Customization & Composability

Retail and eCommerce leaders seeking integration with existing technology stacks without extensive professional services require vendor openness and flexibility. The heatmap reveals notable gaps among 16 vendors, particularly concerning Upload Own Models and Direct SQL Access.

[\[Explore the complete Customization & Composability capability map on page 19\]](#)

In-Session Activation

Real-time, personalized activation capabilities are essential for immediate customer interactions such as mid-session personalization and Buy Online Pickup In-Store (BOPIS). The heatmap demonstrates strong alignment from approximately half the vendors, with the remaining half lacking in one or two critical areas, including Unified Event Collection, Fast Inventory Sync, Sub-Second Decision APIs, Audience Syndication, and Real-Time Profile API.

[\[Explore the full In-Session Activation capability map on page 25\]](#)

Now that we have discussed high-level market capability coverage, let's dive into each archetype and capability lens to examine in-depth vendor mappings and alignments.



Check your preferred vendor's ranking

Click on your vendor to jump straight to its detail page.

* If your PDF viewer disables links, use the page numbers of each vendor section.

Independent CDP (Marketer-led)		Customer Data Infrastructure / Composable		Engagement Hub	
36	Contentstack	82	Netcore	123	Braze
28	Redpoint	74	RudderStack	107	Dotdigital
40	Session AI	78	Tealium	111	Listrak
32	Uniphore	86	Treasure Data	119	Microsoft Dynamics 365 Customer Insights CDP
45	Amperity	90	Twilio Segment	103	Ometria
Commerce & Personalization Engine		Suite CDP		Zeta Global	
65	Bloomreach	145	mParticle	115	Oracle Unity CDP
55	Iterable	157	ReactorData	131	Blueshift
61	SAP Emarsys	149	Salesforce Data 360	135	Marigold
69	Algonomy	153	Adobe Experience Platform	140	Optimove

Vendor Maps: Archetypes and Capability Lenses

In this chapter, we provide retail and eCommerce leaders with a detailed framework for evaluating vendor capabilities aligned closely with their strategic needs and priorities. First, we explore vendor archetypes and individual vendor performance on pages 4-8, offering detailed vendor maps that illustrate the strengths and positioning of each vendor. Subsequently, on pages 9-15, we analyze vendors through four specific Capability Lenses, addressing critical strategic focus areas in depth.

Before we dive deeper, let's clarify key definitions and terms used throughout this analysis:

Understanding the VOX Score

The VOX Score is our proprietary measure reflecting leader priorities across 16 key capability questions. Each question is weighted based on importance and leader consensus. A vendor's total VOX Score sums the weights for each capability they address natively.

Preference Coverage Explained

Preference Coverage counts how many of the 16 critical capabilities each vendor positively addresses, highlighting the breadth of their capabilities.

Context on 16 Capability Questions

Capabilities evaluated include:

- Real-time data observability and lineage tracking
- Comprehensive native integration across major commerce and service platforms
- Unified multi-channel event collection
- Deterministic identity resolution for unified customer profiles
- Near-instant synchronization of inventory and personalization decisions
- Predictive merchandising and real-time personalization capabilities
- Custom model deployment and direct SQL analytics
- Data quality management and automatic anomaly detection
- Field-level role-based access and immutable audit trails
- Instant compliance with privacy and consent regulations
- Regional data residency and encryption
- Real-time audience segmentation and syndication
- Transparent, usage-based pricing



Vendor Archetypes

We've categorized vendors into five mutually exclusive archetypes for clarity:

- **Independent CDP (Marketer-led):** Specialized CDPs for data unification, segmentation, and activation.
- **Customer Data Infrastructure/Composable:** Advanced data infrastructure solutions with pipelines, governance, and warehouse-native capabilities.
- **Engagement Hub:** Comprehensive platforms offering integrated journey orchestration and messaging.
- **Commerce & Personalization Engines:** Solutions optimized for real-time decision-making and merchandising.
- **Suite CDP:** Integrated within broader marketing or CRM suites.

What Are the Color Tags For?

To quickly identify vendor strengths, VOX uses color-coded tags that clearly indicate key capabilities relevant to retail and eCommerce leader priorities:

- **Retail Specialist:** Optimized specifically for retail and e-commerce merchandising scenarios, including product catalog awareness, PIM/OMS integrations, and commerce-grade personalization.
- **Real-time/Edge:** Engineered to act within live customer sessions, providing sub-second decisions, streaming events, and edge runtimes that personalize experiences before customers leave the page.
- **Engagement Hub:** Offers built-in journey orchestration and cross-channel messaging capabilities (email, SMS, push notifications, on-site), integrated with robust customer profiles to support comprehensive campaign management.
- **Warehouse-native/Composable:** Designed for easy integration with your existing data stack, featuring direct SQL or lake-house access, open APIs and connectors, and support for deploying custom analytics models, allowing teams to fully leverage existing technologies.
- **Governance & Residency:** Combines strong governance features like field-level role-based access controls (RBAC) and immutable audit logs with advanced data-sovereignty controls, including options for regional data residency and encryption, suitable for EU-only or US-only deployments.
- **Pricing Transparency:** Clearly published, usage-based pricing independent of professional services, making costs predictable and transparent based on events, monthly tracked users (MTUs), and storage usage.
- **Governance:** Provides robust governance through detailed access controls and comprehensive auditability, including field-level RBAC, immutable logging, and consent enforcement. (Note: Unlike "Governance & Residency," this does not include regional residency capabilities.)
- **Loyalty:** Includes native loyalty management features such as status tracking, points management, personalized offers, and loyalty-driven segmentation directly within customer profiles.



Capability Lens Explained

Capability Lenses offer focused insights into critical areas of interest, though these lenses are not mutually exclusive. We explore four primary lenses, covering all 31 vendors:

- Data Visibility & Quality
- Security, Consent & Residency
- Customization & Composability
- In-Session Activation

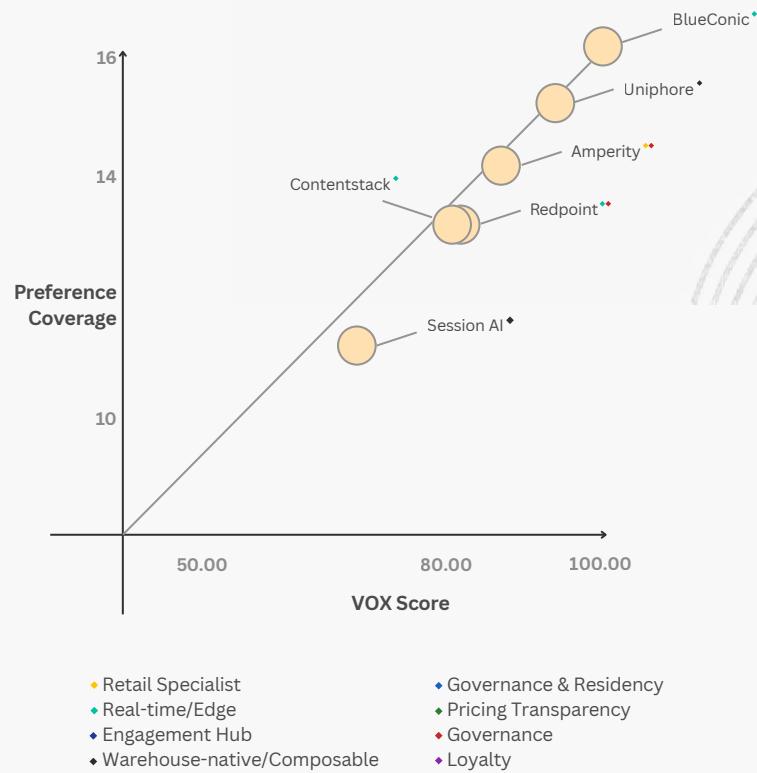
IMPORTANT NOTE

These insights do not designate winners but illustrate vendor alignment with critical capabilities valued by retail and eCommerce leaders. Use these insights as a starting point to shortlist vendors and further investigate their capabilities as needed.

Archetype Overview: Independent CDPs

This map highlights how independent, marketer-led Customer Data Platforms (CDPs) align with the key priorities of enterprise retail and e-commerce leaders. **It evaluates CDPs based on two dimensions—VOX Score (weighted alignment with leader priorities) and Preference Coverage (breadth of capabilities).**

Vendor	VOX Score	Preference Coverage
BlueConic	100	16
Uniphore	93.91	15
Amperity	87.22	14
Redpoint	81.59	13
Contentstack	81.06	13
Session AI	69.14	11



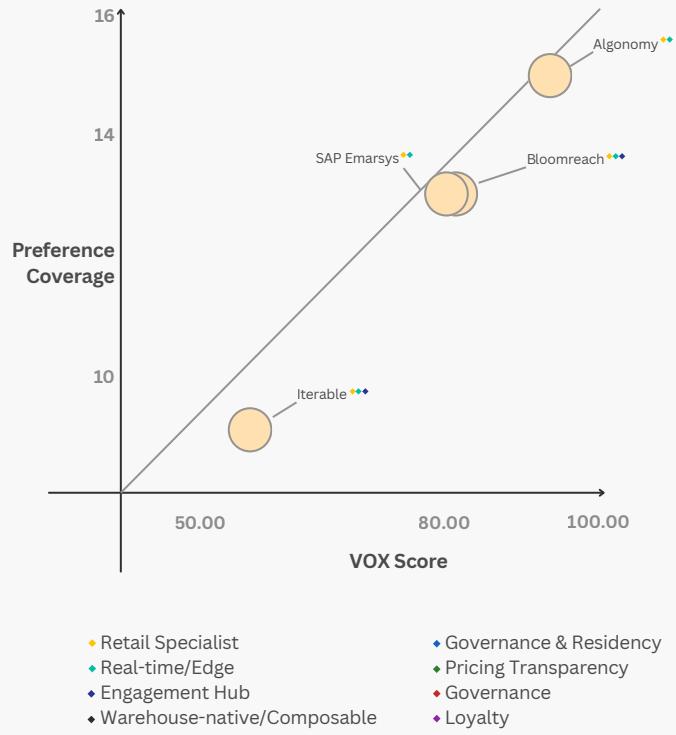
Key Insights

- All featured vendors address at least 10 of the 16 critical capabilities.
- BlueConic, Uniphore, and Amperity notably excel, each meeting 14 or more capabilities with high VOX Scores (80 or higher), signifying strong alignment with leadership priorities.
- Contentstack and Redpoint Global offer balanced capability coverage and solid VOX Scores, indicating well-rounded performance.
- All highlighted vendors demonstrate at least two key strengths, notably focusing heavily on **Real-time/Edge** capabilities (50% coverage), ideal for immediate in-session actions.
- Strong coverage of **Warehouse-native/Composable** and **Governance** features (33% each) supports flexible integrations with existing data infrastructure.
- Select vendors such as BlueConic, Uniphore, and Amperity notably align with specialized retail needs, combining real-time responsiveness and technical openness.

Archetype Overview: Commerce & Personalization Engine

This map compares commerce-first personalization engines integrated with CDP capabilities to drive on-site and in-app personalization. Vendors are evaluated across **two key dimensions**—VOX Score (alignment with enterprise retail and e-commerce priorities) and Preference Coverage (breadth of capabilities).

Vendor	VOX Score	Preference Coverage
Algonomy	93.91	15
Bloomreach	81.12	13
Iterable	56.16	9
SAP Emarsys	81.06	13



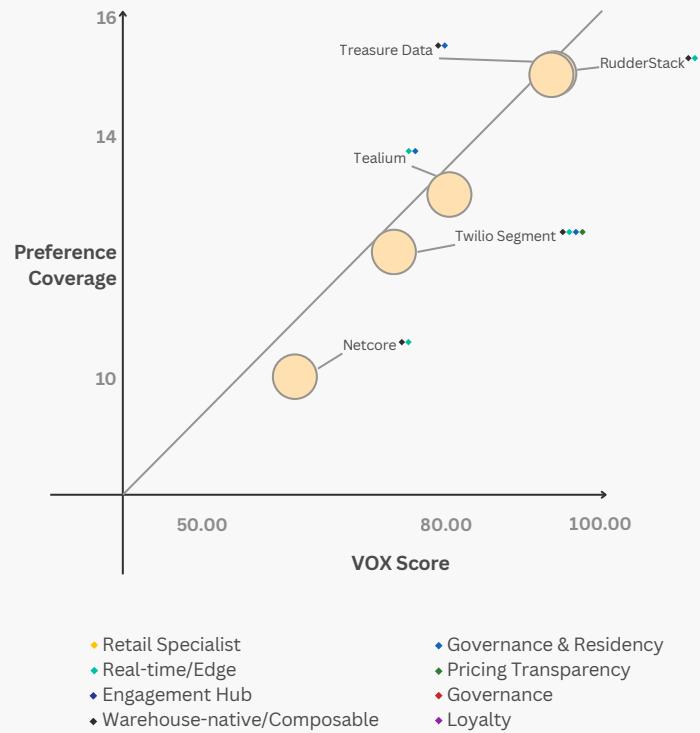
Key Insights

- Algonomy emerges prominently, covering 15 of 16 capabilities with a VOX Score of approximately 93.9, reflecting strong, broad capabilities, particularly for retail merchandising and real-time decision-making.
- Bloomreach and SAP Emarsys demonstrate balanced performance, each covering 13 capabilities with VOX Scores around 81–82, underscoring their robust capabilities in data collection, activation, and commerce-focused personalization.
- Iterable occupies a distinct position, reflecting its strong roots in journey orchestration, with coverage of 9 capabilities and a VOX Score around 56; it is shortlisted by leaders prioritizing comprehensive journey-based personalization.
- All highlighted vendors provide strong and uniform coverage of both Retail Specialist and Real-time/Edge capabilities, ideal for merchandising and immediate in-session personalization.
- Half also include native Engagement Hub features, supporting deeper journey orchestration and multi-channel activation.

Archetype Overview: Customer Data Infrastructure / Composable

This map evaluates engineering-led customer data infrastructure (CDI) solutions designed for building composable Customer Data Platforms (CDPs) within your existing technology stack. Vendors are assessed based on **two key dimensions**—VOX Score (weighted alignment with enterprise leader priorities) and Preference Coverage (breadth of capabilities).

Vendor	VOX Score	Preference Coverage
Netcore	62.32	10
RudderStack	93.98	15
Tealium	81.06	13
Treasure Data	93.91	15
Twilio Segment	74.3	12



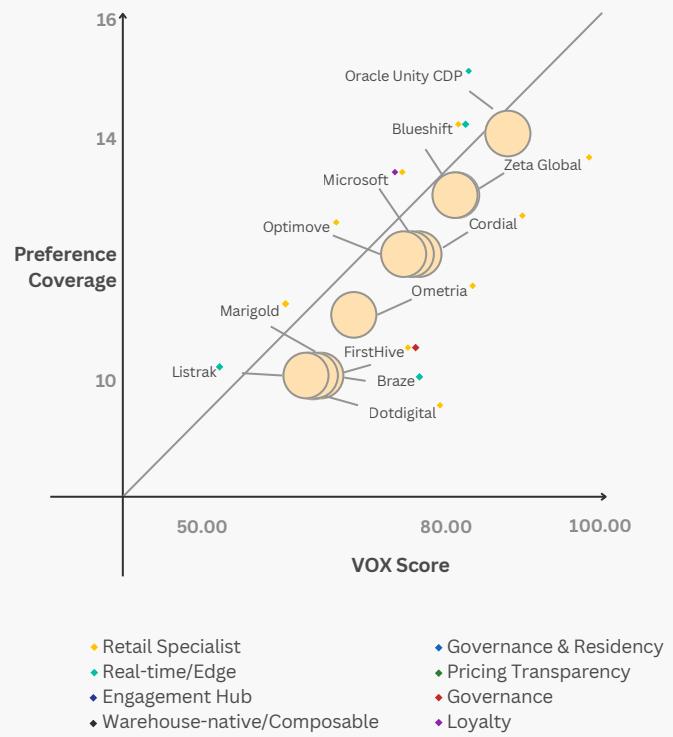
Key Insights

- All vendors highlighted here cover at least 10 of the 16 critical capabilities.
- RudderStack and Treasure Data clearly stand out, each covering 14 or more capabilities with VOX Scores exceeding 80, demonstrating strong alignment with key enterprise priorities.
- Tealium excels particularly in composable and governance, covering 13 capabilities with a strong VOX Score.
- Twilio Segment offers comprehensive pipelines and extensive integrations, demonstrating strong composable infrastructure capabilities.
- Vendors strongly emphasize both **Warehouse-native/Composable** and **Real-time/Edge** capabilities (each at 80%), making them ideal for teams prioritizing architectural flexibility and real-time responsiveness.
- A majority (60%) also deliver substantial Governance & Residency controls, balancing technical agility with compliance needs.

Archetype Overview: Engagement Hub

This map highlights how independent, marketer-led Customer Data Platforms (CDPs) align with the key priorities of enterprise retail and e-commerce leaders. **It evaluates CDPs based on two dimensions—VOX Score (weighted alignment with leader priorities) and Preference Coverage (breadth of capabilities).**

Vendor	VOX Score	Preference Coverage
Blueshift	81.06	13
Braze	62.85	10
Cordial	75.37	12
Dotdigital	62.45	10
FirstHive	62.72	10
Listrak	62.32	10
Marigold	62.38	10
Microsoft Dynamics 365 Customer Insights CDP	75.17	12
Ometria	68.41	11
Optimove	74.77	12
Oracle Unity CDP	87.75	14
Zeta Global	81.06	13



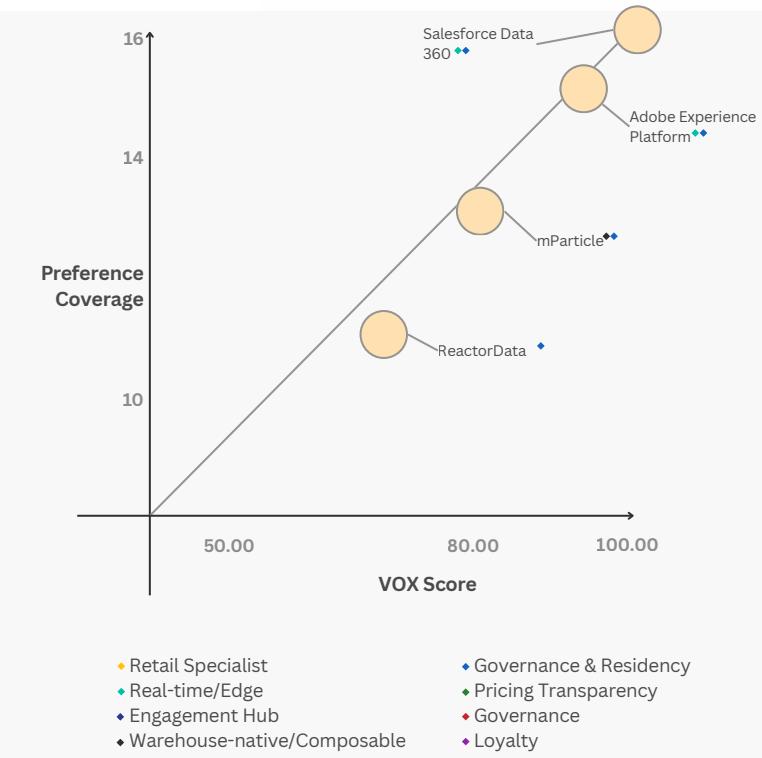
Key Insights

- All vendors displayed cover at least 10 critical capabilities.
- Oracle Unity CDP clearly excels, meeting 14 or more capabilities and achieving a high VOX Score (≥ 80), commonly selected alongside Oracle CX for comprehensive customer engagement.
- Blueshift and Zeta Global also stand out, each with strong VOX Scores (≥ 80) and covering 13 capabilities, highlighting their robust activation capabilities.
- Vendors consistently demonstrate a strong **Retail Specialist** orientation (75%), effectively supporting retail-specific orchestration and campaign management.
- Robust **Real-time/Edge** features emerge notably in one-third of vendors, enabling mid-session personalization.

Archetype Overview: Suite CDP

This map highlights how independent, marketer-led Customer Data Platforms (CDPs) align with the key priorities of enterprise retail and e-commerce leaders. **It evaluates CDPs based on two dimensions—VOX Score (weighted alignment with leader priorities) and Preference Coverage (breadth of capabilities).**

Vendor	VOX Score	Preference Coverage
Adobe Experience	93.91	15
mParticle	80.86	13
ReactorData	69.01	11
Salesforce Data 360	100	16



Key Insights

- All featured vendors cover at least 11 critical capabilities.
- Salesforce Data 360 and Adobe Experience Platform are clear leaders, covering 14 or more capabilities each, with VOX Scores above 80.
- mParticle achieves a strong VOX Score (≥ 80) covering 13 capabilities, notably effective for data capture and governance within broader suite ecosystems.
- All featured vendors strongly prioritize Governance & Residency features, providing comprehensiveness, auditability, compliance, and data sovereignty.
- Half of the vendors also deliver notable Real-time/Edge capabilities, supporting in-session customer interactions.

Capability Lens - Data Visibility & Quality

Vendor	Live Data Lineage	Direct SQL Access	Anomaly Quarantine
Adobe Experience Platform			
Amperity			
BlueConic			
Contentstack			
mParticle			
ReactorData			
Redpoint			
RudderStack			
Salesforce Data 360			
Session AI			
Tealium			
Treasure Data			
Twilio Segment			
Uniphore			

This capability lens addresses a crucial buyer question: "Can I clearly observe and reliably trust my data paths and health?" It evaluates vendor capabilities related to real-time visibility, data access, and proactive quality management.

- Live Data Lineage: Capability to deliver dashboards with refresh intervals ≤ 60 seconds (p95) showcasing complete source-to-destination data lineage.
- Direct SQL Access: Provision of a direct SQL workspace or lake-house access for flexible data analysis.
- Anomaly Quarantine: Continuous monitoring for data quality with automated quarantine and alerts for detected anomalies.

Key Insights

- 14 of 31 vendors address the Data Visibility & Quality capability lens, covering at least one of the evaluated areas — Live Data Lineage, Direct SQL Access, or Anomaly Quarantine.
- The majority of these 14 vendors demonstrate full coverage across all three capabilities, indicating strong overall maturity in data transparency and governance.
- Only a small subset of vendors show partial coverage, typically lacking automated anomaly quarantine or complete lineage visibility — areas that still distinguish the most advanced data infrastructures in this cycle.

Capability Lens - Security, Consent & Residency

Vendor	RBAC & Audit	Instant Consent Suppression	Residency & Encryption
Adobe Experience Platform			
Amperity			
BlueConic			
Contentstack			
mParticle			
Netcore			
ReactorData			
Redpoint			
RudderStack			
Salesforce Data 360			
Session AI			
Tealium			
Treasure Data			
Twilio Segment			
Uniphore			

This capability lens evaluates vendor solutions addressing critical questions around security, compliance, and data governance, focusing specifically on roles, permissions, consent management, and data residency.

- RBAC & Audit: Vendors providing robust, field-level role-based access control and immutable audit logging.
- Instant Consent Suppression: Built-in consent and preference management systems that immediately halt data activation upon permission changes.
- Residency & Encryption: Solutions offering regional data residency compliance options with robust encryption at rest and in transit.

Key Insights

- 15 of 31 vendors address the Security, Consent & Residency capability lens.
- The majority of these **15 vendors demonstrate comprehensive coverage across all three evaluated capabilities**—RBAC & Audit, Instant Consent Suppression, and Residency & Encryption—reflecting strong alignment with enterprise governance and compliance expectations.
- Only a small subset of vendors show partial coverage, often lacking Instant Consent Suppression, an area that remains a notable differentiator among leading compliance-focused solutions.

Capability Lens - Customization & Composability

Vendor	Native Inbound Connectors	Embedded Predictive Models	Upload Own Models	Direct SQL Access	Real-Time Profile API
Adobe Experience Platform	■	■	■	■	■
Amperity	■	■	■	■	■
BlueConic	■	■	■	■	■
Contentstack	■	■	■	■	■
mParticle	■	■	■	■	■
Netcore	■	■	■	■	■
ReactorData	■	■	■	■	■
Redpoint	■	■	■	■	■
RudderStack	■	■	■	■	■
Salesforce Data 360	■	■	■	■	■
Session AI	■	■	■	■	■
Tealium	■	■	■	■	■
Treasure Data	■	■	■	■	■
Twilio Segment	■	■	■	■	■
Uniphore	■	■	■	■	■

This capability lens assesses vendor flexibility and openness, particularly around customization and composability, answering the critical question: "Can I seamlessly leverage my existing models and technology stack without extensive professional services?"

- **Embedded Predictive Models:** Availability of predictive models easily deployed by marketers.
- **Upload Own Models:** Support for user-generated machine learning models (PMML, ONNX, Python).
- **Direct SQL Access:** Capability to directly interact with your existing lake-house infrastructure via SQL.
- **Real-Time Profile API:** Access to unified customer profiles via real-time queries or streams.
- **Native Inbound Connectors:** Out-of-the-box connectors for major commerce, marketing, and service platforms.

Key Insights

- Only **9 vendors** comprehensively cover all five capabilities, providing greater flexibility and reduced reliance on external services.
- Remaining vendors show gaps, especially in **Upload Own Models** and **Direct SQL Access**, important for composable, custom integrations.

Capability Lens - In-Session Activation

Vendor	Unified Event Collection	Fast Inventory Sync	Sub-Second Decision APIs	Audience Syndication	Real-Time Profile API
Algonomy					
Bloomreach					
Blueshift					
Braze					
Cordial					
Dotdigital					
FirstHive					
Iterable					
Listrak					
Marigold					
Microsoft Dynamics 365					
Ometria					
Optimove					
Oracle Unity CDP					
SAP Emarsys					
Zeta Global					

This capability lens evaluates vendor solutions based on their ability to perform personalized activations in real-time, directly while the customer interacts on-site or in-app.

- Unified Event Collection (Q3): Capability for server-side and client-side data collection across web, apps, email, and POS, consolidated into one SDK.
- Fast Inventory (Q5): Synchronization of online and in-store inventory within less than 5 minutes (p95).
- Sub-Second Decision (Q6): Availability of APIs for decision-making and personalization, responding in under 1 second (p95).
- Audience Syndication (Q14): Native capabilities for reverse-ETL and integration with various audience destinations.
- Real-Time Profile API (Q15): Ability to query and subscribe to unified customer profiles in real-time.

Key Insights

- Only **8 vendors** comprehensively cover all five capabilities, fully enabling real-time, in-session personalization and activation.
- Remaining vendors have gaps, primarily in **Fast Inventory Sync** and **Audience Syndication**, important for effective real-time customer interactions and session conversion.

Vendor Deep Dives

Click on your vendor to jump straight to its detail page.

* If your PDF viewer disables links, use the page numbers of each vendor section.

Independent CDP (Marketer-led)	Customer Data Infrastructure / Composable	Engagement Hub
36 Contentstack	82 Netcore	123 Braze
28 Redpoint	74 RudderStack	107 Dotdigital
40 Session AI	78 Tealium	95 FirstHive
32 Uniphore	86 Treasure Data	111 Listrak
45 Amperity	90 Twilio Segment	119 Microsoft Dynamics 365 Customer Insights CDP
		103 Ometria
		127 Zeta Global
Commerce & Personalization Engine	Suite CDP	99 Cordial
65 Bloomreach	145 mParticle	115 Oracle Unity CDP
55 Iterable	157 ReactorData	131 Blueshift
61 SAP Emarsys	149 Salesforce Data 360	135 Marigold
69 Algonomy	153 Adobe Experience Platform	140 Optimove

Independent CDPs (Marketer-led)



Vendor Deep Dive



Redpoint

Primary Archetype: Independent CDP (Marketer-led)

Strengths: Real-time/Edge; Governance

Capability Lens: Data Visibility & Quality; Security, Consent & Residency; Customization & Composability

VOX Score: 81.59 / 100

Preference Coverage: 13 / 16 questions “Yes”

📞 Talk to Redpoint: <https://redpointglobal.com/>

📞 [Ask VOX for Introduction](#)

Oct 2025 CDP Pulse Report - Redpoint

Yes/No	Details	Source(s)	
1. Does your solution provide data-observability dashboards that refresh in ≤ 60 seconds (p95) and show source-to-destination lineage for every first-party data stream?	Yes	Redpoint includes real-time data observability dashboards. Users get up-to-the-moment visibility into data pipelines, with full lineage from each source to its destination, ensuring every first-party data stream is tracked end-to-end.	https://www.redpointglobal.com/data-observability/
2. Does your solution ship native inbound connectors or APIs for the top-10 commerce, marketing-automation, and service platforms without custom code or PS?	Yes	Redpoint offers a broad range of out-of-the-box connectors and APIs. It can natively ingest data from major commerce, marketing automation, and service platforms (CRM, CMS, email/SMS providers, ad platforms, etc.) and activate data to them, all without the need for custom coding or professional services.	https://www.redpointglobal.com/integrations/
3. Does your solution support server-side and client-side collection for web, mobile-app, email, and in-store (POS/scanner) events in a single SDK?	No	While Redpoint can collect and unify data from web, mobile, email, and in-store sources, it does not provide a single integrated SDK covering both client-side and server-side event capture across all channels. Data collection is achieved via connectors, tag management, and APIs, but public documentation does not indicate one unified SDK for all event types.	
4. Does your solution offer deterministic identity-resolution that stitches online cookies, MAIDs, loyalty IDs and POS receipts into one customer profile?	Yes	Redpoint provides powerful identity resolution that unifies online and offline identifiers into a single customer profile. It links cookies and device IDs with loyalty program IDs and in-store/POS data, using primarily deterministic matching (augmented by probabilistic techniques) to stitch together a complete view of each individual across all touchpoints.	https://www.redpointglobal.com/identity-resolution/ https://www.redpointglobal.com/blog/what-is-identity-resolution/
5. Does your solution sync online and in-store inventory events in < 5 minutes (p95) to enable BOPIS and endless-aisle use cases?	Yes	Redpoint's platform supports near-real-time data integration between online and in-store systems. It can update and unify inventory events from e-commerce and POS within minutes, enabling use cases like "buy online, pick up in store" and endless aisle. Redpoint emphasizes real-time inventory visibility so that stores have up-to-date knowledge of online orders and inventory, ensuring seamless BOPIS experiences.	https://www.redpointglobal.com/blog/if-your-cx-strategy-doesnt-include-bopis-youre-doing-it-wrong/ https://www.redpointglobal.com/blog/the-power-of-the-golden-record-for-data-driven-retailers/
6. Does your solution expose decision or personalization APIs that return in < 1 second (p95) so actions fire while the session is active?	Yes	Redpoint offers real-time decisioning APIs with sub-second response performance. The platform's personalization engine can evaluate context and return next-best actions or content in milliseconds, allowing personalized offers or recommendations to be delivered immediately during an active user session.	businesswire.com/news/home/2023052005057/en/Redpoints-Leading-CDP-Processes-100s-of-Billions-of-Interactions-a-Year-for-Some-of-the-Worlds-Biggest-Companies
7. Does your solution include embedded predictive-merchandising models (e.g., next-best product) deployable by business users without data-science help?	Yes	Redpoint's CDP comes with built-in automated machine learning models for marketing. Non-technical users can leverage features like next-best-action and product recommendation models directly in the platform. These predictive merchandising models are provided out-of-the-box and can be used by marketers to personalize experiences (e.g. suggesting the next best product) without needing a data scientist to build them.	https://www.redpointglobal.com/wp-content/uploads/2020/05/Redpoint-Automated-Machine-Learning.pdf

Oct 2025 CDP Pulse Report - Redpoint

Yes/No	Details	Source(s)
8. Does your solution allow teams to upload and deploy their own ML/AI models (PMML/ONNX/Python) to score profiles or events in real time?	Yes Redpoint supports custom analytics model integration. Teams have the flexibility to configure and add their own machine learning models into the platform to score customers or events. This means organizations can bring proprietary models (e.g. built in Python or exported to PMML/ONNX) and deploy them within Redpoint's real-time decisioning workflows to tailor scoring and predictions to their business.	https://www.cdpinstitute.org/wp-content/uploads/2021/04/Redpoint-2172-Real-Time-Customer-Engagement.pdf
9. Does your solution provide an SQL workspace or direct lake-house access for segmentation, attribution, and ad-hoc analytics?	Yes Redpoint enables flexible, SQL-based access to customer data. The platform provides a SQL workspace for advanced segmentation and analysis (users can write SQL queries to define audience segments), and it natively connects with cloud data warehouses/data lakes (Snowflake, BigQuery, etc.) to allow direct querying and analytics on customer data. This ensures teams can perform ad-hoc analysis and attribution using familiar SQL tools or BI integrations.	https://docs.redpointglobal.com/cdp/compose-sql-segment https://business.adobe.com/content/dam/dx/uk/en/resources/reports/pdf/MIR_CDP_4th_Adobe_uk.pdf
10. Does your solution run continuous data-quality and anomaly detection with automatic quarantine or alerting before bad data feeds downstream systems?	Yes Redpoint continuously monitors incoming data quality in real time. It performs automated validation checks and anomaly detection on data as it is ingested, catching errors or inconsistencies early. The platform provides timely alerts/visibility into data issues and can stop or flag "bad" data before it propagates to downstream systems, thereby protecting data consumers from quality problems.	https://www.redpointglobal.com/data-observability/ https://www.redpointglobal.com/data-observability/
11. Does your solution enforce role-based access controls and immutable audit logs at the field level across all regions?	Yes Redpoint includes robust data governance features such as role-based access control (RBAC) and detailed auditing. User roles and permissions can be defined to the field/data element level, and all user activities are recorded in auditable logs. These controls are consistent across regions, supporting compliance and security requirements globally (with immutable logs for traceability).	rudderstack.com/competitors/rudderstack-cdp-vs-redpoint-global-cdp/
12. Does your solution offer built-in consent & preference management that suppresses activation instantly when permissions change (GDPR/CCPA ready)?	No Redpoint does not natively provide a full consent/preferences management module. Instead, it is designed to integrate with dedicated consent management systems (e.g. PossibleNow or other preference center tools). While Redpoint will honor consent flags in data (and can immediately suppress outreach to customers who opt out), the solution relies on external consent management platforms to collect and manage those preferences in compliance with GDPR/CCPA, rather than offering its own built-in preference center.	
13. Does your solution support regional data residency options (e.g., EU-only, US-only) with encryption at rest (AES-256) and in transit (TLS 1.2+)?	Yes Redpoint can be deployed to meet data residency requirements. It offers cloud deployment options in multiple regions (for example, in North America or Europe) to keep customer data in-region as needed. The platform also adheres to strict security standards: customer data is encrypted at rest and in transit, aligning with industry best practices (AES-256 encryption for stored data and TLS 1.2+ for data in motion, among the measures under its ISO 27001 and SOC 2 compliant security program).	https://docs.redpointglobal.com/trust/redpoint-s-security-standards-and-compliance https://docs.redpointglobal.com/trust/redpoint-s-security-standards-and-compliance
14. Does your solution provide reverse-ETL / audience-syndication connectors that push profiles and events to major ad, analytics, and messaging destinations natively?	Yes Redpoint supports "reverse ETL" style activation through its many native outbound connectors. The platform can syndicate customer profiles and events directly to a wide array of downstream marketing and analytics destinations (advertising networks, social platforms, email/SMS delivery systems, web personalization tools, analytics services, etc.). These native connectors (100+ out-of-the-box) let Redpoint push audience data or event streams into external systems without custom development.	https://business.adobe.com/content/dam/dx/uk/en/resources/reports/pdf/MIR_CDP_4th_Adobe_uk.pdf https://www.redpointglobal.com/integrations/

Oct 2025 CDP Pulse Report - Redpoint

Yes/No	Details	Source(s)	
15. Does your solution expose a unified customer-profile API/stream that any downstream system can query or subscribe to in real time?	Yes	<p>Redpoint provides real-time access to the unified customer profile via open APIs and streaming integrations. It offers a comprehensive REST API through which external systems can query customer profiles or trigger interactions on demand. Additionally, Redpoint can stream data (e.g., via Kafka or similar) to let downstream applications subscribe to profile updates or events as they happen, ensuring up-to-date information is available in real time.</p>	https://business.adobe.com/content/dam/dx/uk/en/resources/reports/pdf/MIR_CDP_4th_Adobe_uk.pdf https://www.redpointglobal.com/integrations/
16. Does your solution publish transparent, usage-based pricing (events, MTU, storage) in writing, independent of professional-services spend?	No	<p>Redpoint's pricing model is not publicly disclosed as a transparent usage-based rate card. The company typically negotiates pricing based on factors like number of customer records and data volume, and it does not publish fixed per-event or per-user pricing on its website. In practice, customers must engage with Redpoint for a tailored quote; pricing is not provided in writing independent of services.</p>	

Vendor Deep Dive



Uniphore

Primary Archetype: Independent CDP (Marketer-led)

Strengths: Warehouse-native/Composable

Capability Lens: Data Visibility & Quality; Security, Consent & Residency; Customization & Composability

VOX Score: 93.91 / 100

Preference Coverage: 15 / 16 questions “Yes”

📞 Talk to Uniphore: <https://www.uniphore.com/>

📞 [Ask VOX for Introduction](#)

Oct 2025 CDP Pulse Report - Uniphore

Yes/No	Details	Source(s)
1. Does your solution provide data-observability dashboards that refresh in \leq 60 seconds (p95) and show source-to-destination lineage for every first-party data stream?	Yes The platform includes dashboards to monitor data ingestion pipelines in real time. Users can view and control data pipelines via a dashboard, which provides visibility into data flows. Furthermore, the solution's data layer (augmented by Infoworks technology) automatically discovers and catalogs enterprise data, helping track data from source to destination.	https://www.actioniq.com/wp-content/uploads/CXHub_SolutionBrief_E.pdf
2. Does your solution ship native inbound connectors or APIs for the top-10 commerce, marketing-automation, and service platforms without custom code or PS?	Yes Uniphore's CDP (ActionIQ) offers out-of-the-box connectors and APIs for a wide range of systems across commerce, marketing automation, CRM, service, etc. The integration library covers major platforms (e.g. Adobe, Salesforce, e-commerce and advertising tools) so data can be ingested without custom development or professional services.	https://www.uniphore.com/products/marketing-ai/cdp-agent/integrations/ https://www.uniphore.com/products/marketing-ai/cdp-agent/integrations/
3. Does your solution support server-side and client-side collection for web, mobile-app, email, and in-store (POS/scanner) events in a single SDK?	Yes The solution can capture events from both client-side (e.g. via an ActionIQ Tag on websites) and server-side sources (via APIs). It supports real-time event streams such as website page views, mobile app events, email opens, and even in-store/POS transactions. These are collected through a unified framework (streaming API/SDK), allowing multiple channels' events to feed into the CDP's single customer view.	https://hightouch.com/compare-cdps/actioniq-vs-mparticle https://hightouch.com/compare-cdps/actioniq-vs-mparticle
4. Does your solution offer deterministic identity-resolution that stitches online cookies, MAIDs, loyalty IDs and POS receipts into one customer profile?	Yes The platform includes a robust identity resolution system. It deterministically links identifiers across anonymous and known sources (e.g. web cookies, mobile device IDs, loyalty program IDs, in-store purchase records) to build a unified customer profile. This identity framework can join first-party identifiers and also leverage third-party identity services, ensuring one profile per customer.	https://info.actioniq.com/hubfs/AIQ%20Identity%20Solution%20Brief%20(1).pdf https://info.actioniq.com/hubfs/AIQ%20Identity%20Solution%20Brief%20(1).pdf
5. Does your solution sync online and in-store inventory events in < 5 minutes (p95) to enable BOPIS and endless-aisle use cases?	Yes The CDP supports near-real-time data streaming, so inventory updates from online or in-store systems can be ingested and available within minutes. It combines streaming data with historical data to inform real-time experiences. This low-latency data sync enables use cases like Buy Online Pickup In-Store and endless aisle, as inventory changes are reflected promptly (often within a few minutes) across channels.	https://www.uniphore.com/products/marketing-ai/cdp-agent/ https://actioniq.magnetikwebdev.com/solutions/
6. Does your solution expose decision or personalization APIs that return in < 1 second (p95) so actions fire while the session is active?	Yes The solution provides real-time decisioning APIs – notably a Customer 360 “Profile API” – that external systems can call to get customer attributes or segment membership in milliseconds. These APIs are designed for in-session personalization (e.g. web personalization engines or next-best-action systems can fetch a profile and get a response fast, usually well under one second, so they can react during an active user session).	https://info.actioniq.com/hubfs/Product%20Intelligence/AIQ%20RTCX%20Solution%20Brief.pdf https://info.actioniq.com/hubfs/Product%20Intelligence/AIQ%20RTCX%20Solution%20Brief.pdf
7. Does your solution include embedded predictive-merchandising models (e.g., next-best product) deployable by business users without data-science help?	Yes Uniphore's CDP comes with a library of pre-built predictive models that non-technical users can leverage. These native models (for common use cases like churn likelihood, lifetime value, next-best action/product recommendations, optimal send time, etc.) are available out-of-the-box. They can be configured with business rules and deployed by marketers directly through the platform's UI (no data science coding needed).	https://info.actioniq.com/hubfs/SolutionBrief https://info.actioniq.com/hubfs/SolutionBrief

Oct 2025 CDP Pulse Report - Uniphore

Yes/No	Details	Source(s)
8. Does your solution allow teams to upload and deploy their own ML/AI models (PMML/ONNX/Python) to score profiles or events in real time?	Yes The platform supports bringing your own models. Data science teams can integrate custom models (built externally in Python/ML frameworks, exported e.g. via PMML/ONNX) into the CDP's machine learning framework. These "hosted models" can then be operationalized within the CDP – meaning the model will run scoring on customer profiles or incoming events in real-time as part of the platform's workflows.	https://info.actioniq.com/hubfs/SolutionBrief%20ActionIQ/FY2021%20Q3%20Feb19%20Solution%20Brief.pdf https://info.actioniq.com/hubfs/SolutionBrief%20ActionIQ/FY2021%20Q3%20Feb19%20Solution%20Brief.pdf
9. Does your solution provide an SQL workspace or direct lake-house access for segmentation, attribution, and ad-hoc analytics?	Yes The solution offers a self-service SQL workspace (Query Service) for advanced users. Technical teams can write and execute SQL queries on data within the CDP or even directly on connected data warehouses/lakehouses thanks to the composable architecture. This empowers users to do custom segmentation, attribution analysis, and other ad-hoc analytics by querying the unified customer data with SQL.	https://www.actioniq.com/blog/query-service/ https://www.actioniq.com/blog/looking-back-2023/
10. Does your solution run continuous data-quality and anomaly detection with automatic quarantine or alerting before bad data feeds downstream systems?	Yes The platform includes data quality monitoring capabilities. Users can define variance thresholds and get anomaly alerts if incoming data deviates unexpectedly. This means the system will flag anomalies in data feeds and alert teams so they can stop or quarantine bad data before it propagates to downstream audiences or reports. (Automatic quarantine is handled via configured rules/alerts, ensuring bad data can be caught and addressed proactively.)	https://info.actioniq.com/hubfs/SolutionBrief%20ActionIQ/FY2021%20Q3%20Feb19%20Solution%20Brief.pdf
11. Does your solution enforce role-based access controls and immutable audit logs at the field level across all regions?	Yes The solution provides granular role-based permissions to control who can access or edit data (down to specific fields if needed). It also maintains detailed audit trails of user activity. All changes and accesses are logged, allowing governance teams to see an immutable record of who did what. These security controls apply across regions, supporting global enterprises with strict governance needs.	http://info.actioniq.com
12. Does your solution offer built-in consent & preference management that suppresses activation instantly when permissions change (GDPR/CCPA ready)?	Yes Yes – the platform is designed to comply with GDPR/CCPA and similar regulations. It can store and honor customer consent and preferences (e.g. do-not-contact flags). If a user's consent status changes or they opt-out, the CDP can immediately suppress that individual from any active audiences or campaigns. The system supports privacy-by-design features like fine-grained data purging and respecting "do not sell/share" signals, ensuring that preference changes are reflected instantly in downstream activation.	https://info.actioniq.com/hubfs/AIO%20Identity%20Solution%20Brief%20(1).pdf https://info.actioniq.com/hubfs/ActionIQ%20Industry%20Brief%20Solutions/ActionIQ_Government_Regulations_Compliance_Solution_Brief.pdf
13. Does your solution support regional data residency options (e.g., EU-only, US-only) with encryption at rest (AES-256) and in transit (TLS 1.2+)?	Yes The platform can be deployed in a way that keeps data resident in required regions (cloud, multi-cloud or on-prem as needed to meet data sovereignty). It offers a "zero-copy" architecture that leaves data in-place, helping comply with EU-only or other residency mandates. All data is secured via strong encryption – data at rest is encrypted with AES-256, and all data in transit uses TLS 1.2 or higher.	https://www.uniphore.com/press-releases/uniphore-launches-business-ai-cloud-a-sovereign-composable-secure-ai-platform-to-power-the-agentic-enterprise/ https://www.uniphore.com/security/
14. Does your solution provide reverse-ETL / audience-syndication connectors that push profiles and events to major ad, analytics, and messaging destinations natively?	Yes Yes. The CDP includes native connectors to export and sync customer profiles or event data to a wide array of marketing, advertising, analytics, and messaging platforms. This covers major ad networks, email service providers, CRM systems, business intelligence tools, etc. Marketers can easily syndicate audiences or send event triggers to these downstream systems without custom builds, using the built-in "activation" integrations.	https://www.uniphore.com/products/marketing-ai/cdp-agent/integrations/ https://www.uniphore.com/products/marketing-ai/cdp-agent/integrations/

Oct 2025 CDP Pulse Report - Uniphore

Yes/No	Details	Source(s)	
15. Does your solution expose a unified customer-profile API/stream that any downstream system can query or subscribe to in real time?	Yes	<p>The solution exposes a unified profile API that allows external systems to query a customer's profile on demand in real time. For example, a web or call-center application can request the latest profile data (attributes, segment memberships, etc.) via this API and get a response in milliseconds. This enables downstream systems to subscribe to or fetch customer data in real-time to personalize experiences. (Streaming subscription to profile changes is also supported via real-time event connectors and webhooks.)</p>	<p>https://info.actioniq.com/hubfs/Product%20Intelligence/AIQ%20RTCX%20Solution%20Brief.pdf</p> <p>https://info.actioniq.com/hubfs/Product%20Intelligence/AIQ%20RTCX%20Solution%20Brief.pdf</p>
16. Does your solution publish transparent, usage-based pricing (events, MTU, storage) in writing, independent of professional-services spend?	No	<p>Pricing is not publicly posted. Uniphore/ActionIQ typically provides pricing through custom quotes rather than a transparent rate card. While the pricing model is usage-based (considering factors like event volume, monthly tracked users, data storage, etc.), these details are communicated during the sales process. There is no published pricing page – the cost is tailored per client and not tied to professional services spend.</p>	<p>https://www.g2.com/products/actioniq-by-uniphore/pricing</p>

Vendor Deep Dive



Contentstack

Primary Archetype: Independent CDP (Marketer-led)

Strengths: Real-time/Edge

Capability Lens: Data Visibility & Quality; Security, Consent & Residency; Customization & Composability

VOX Score: 81.06 / 100

Preference Coverage: 13 / 16 questions “Yes”

📞 Talk to Contentstack: <https://www.contentstack.com/>

📞 [Ask VOX for Introduction](#)

Oct 2025 CDP Pulse Report - Contentstack

Yes/No	Details	Source(s)	
1. Does your solution provide data-observability dashboards that refresh in ≤ 60 seconds (p95) and show source-to-destination lineage for every first-party data stream?	No	No. Contentstack CDP (Lytics) does not offer data-observability dashboards with guaranteed p95 refresh rates; its monitoring metrics update at intervals ranging from one minute to one hour. While its UI visualizes data flow between sources and jobs, it does not provide granular, source-to-destination field-level lineage for individual data streams.	https://docs.lytics.com/docs/monitoring-lytics https://docs.lytics.com/docs/dashboard-report https://docs.lytics.com/docs/reports-components
2. Does your solution ship native inbound connectors or APIs for the top-10 commerce, marketing-automation, and service platforms without custom code or PS?	Yes	Yes. Contentstack CDP (Lytics) provides a gallery of pre-built integrations for data ingestion. It natively connects to hundreds of popular martech systems – covering major e-commerce, advertising, CRM, analytics, and email platforms – enabling data import/export without custom code.	https://www.lytics.com/integrations https://docs.lytics.com/docs/integrated-marketing-tools https://docs.lytics.com/docs/lytics-integration-options
3. Does your solution support server-side and client-side collection for web, mobile-app, email, and in-store (POS/scanner) events in a single SDK?	Yes	Yes. Contentstack CDP (Lytics) provides multiple data collection methods that cover client and server. It offers a JavaScript tag for websites, mobile SDKs (iOS/Android) for apps, and open APIs for server-side or offline sources. These allow ingesting events from web, mobile, email, and even point-of-sale systems into one unified pipeline.	https://docs.lytics.com/docs/sdks-web https://docs.lytics.com/docs/sdks-mobile https://docs.lytics.com/docs/lytics-javascript-tag https://docs.lytics.com/docs/lytics-file-service
4. Does your solution offer deterministic identity-resolution that stitches online cookies, MAIDs, loyalty IDs and POS receipts into one customer profile?	Yes	Yes. Contentstack CDP (Lytics) performs deterministic identity resolution by mapping identifiers from various data streams to a unified identity graph using configurable strategies. This process stitches together anonymous and known user data points into a persistent, consolidated customer profile. Lytics reports up to ~99.7% profile matching accuracy across web, email, and offline identifiers.	https://docs.lytics.com/docs/identity-resolution
5. Does your solution sync online and in-store inventory events in < 5 minutes (p95) to enable BOPIS and endless-aisle use cases?	Yes	Yes. The platform's architecture is designed for low-latency event processing, with documented entity updates occurring in "milliseconds to seconds," which is sufficient for real-time use cases. This capability supports the near-instantaneous synchronization of event data, such as inventory changes from a POS system, to enable use cases like BOPIS or endless-aisle. For example, Lytics' Shopify integration brings in real-time order and inventory data (product catalog, stock updates, etc.), allowing the combined platform to reflect inventory changes almost immediately for cross-channel scenarios.	https://docs.lytics.com/docs/data-streams-1 https://docs.lytics.com/docs/shopify https://docs.lytics.com/docs/lytics-and-liveramp
6. Does your solution expose decision or personalization APIs that return in < 1 second (p95) so actions fire while the session is active?	Yes	Yes. Contentstack CDP (Lytics) exposes a low-latency REST API for fetching user profiles and segment memberships in real time, with response times engineered to be under one second. This enables synchronous, in-session personalization by allowing client-side or server-side applications to retrieve user data and trigger an experience while the session is active. The system supports same-page, first-visit personalization for both anonymous and known visitors, using edge delivery for speed. In practice, this means personalization decisions are returned essentially instantly (sub-second) to be applied within the active session, yielding flicker-free, in-session experiences.	https://docs.lytics.com/docs/lytics-platform-data-flow-and-access
7. Does your solution include embedded predictive-merchandising models (e.g., next-best product) deployable by business users without data-science help?	Yes	Yes. Contentstack CDP's (Lytics) embedded AI/ML capabilities focus on building predictive models for user-level behavioral propensities like churn, conversion, or engagement scores. It has codeless tools for predictive segmentation and recommendations – for example, business users can create lookalike propensity models or next-best product recommendations via the UI without coding. Marketers can deploy web product recommendation widgets and build predictive audiences (e.g. "likely to purchase") through Lytics' interface, no data scientist required.	https://docs.lytics.com/docs/getting-started-1

Oct 2025 CDP Pulse Report - Contentstack

Yes/No	Details	Source(s)	
8. Does your solution allow teams to upload and deploy their own ML/AI models (PMML/ONNX/Python) to score profiles or events in real time?	No	No. There is no indication that Contentstack's CDP enables importing custom external ML models. Lytics provides its own machine-learning framework (for lookalike scoring, affinity, etc.) via the UI, but it does not advertise support for uploading PMML/ONNX/Python models for real-time scoring. All ML capabilities appear to be native to the platform's tools. The platform allows users to build custom behavioral and lookalike models using its proprietary interface and machine learning pipeline.	https://docs.lytics.com/docs/building-lookalike-models
9. Does your solution provide an SQL workspace or direct lake-house access for segmentation, attribution, and ad-hoc analytics?	Yes	Yes. Lytics offers direct warehouse access and SQL-based segmentation. Its Cloud Connect feature lets users query data in a cloud data warehouse (e.g. BigQuery) using SQL to build audiences or run analytics. In fact, Contentstack provides an optional managed BigQuery data lake (via "Lytics Conductor") and allows writing SQL queries on profiles/events to create segments, perform attribution, etc., with results usable in the platform.	https://docs.lytics.com/docs/lytics-integration-options https://docs.lytics.com/docs/cloud-connect-intro https://docs.lytics.com/docs/creating-cloud-connect-audiences
10. Does your solution run continuous data-quality and anomaly detection with automatic quarantine or alerting before bad data feeds downstream systems?	No	No. The platform does not feature a continuous, automated system for post-ingestion data quality monitoring, anomaly detection, or the automatic quarantining of anomalous data. Data quality enforcement within the platform is primarily focused on schema validation and error handling at the point of ingestion using LQL. The platform does include monitoring and can send alerts (e.g. email alerts on data pipeline issues or quota thresholds), but it does not mention proactive anomaly detection or auto-quarantining of suspicious data. Data quality management appears to be manual or via configured alerts rather than an automated quarantine system.	https://docs.lytics.com/docs/lytics-monitoring https://docs.lytics.com/docs/data-management https://docs.lytics.com/docs/background-d-processing
11. Does your solution enforce role-based access controls and immutable audit logs at the field level across all regions?	Yes	Yes. Contentstack CDP (Lytics) enforces role-based access controls (RBAC) that can be configured to restrict user permissions at the UI and API level, including field-level restrictions for PII (e.g., marking certain fields as private/PII so only authorized roles can view them). The platform also maintains audit logs of content and schema changes (available in higher-tier plans) to track user actions. These controls apply across its regional deployments.	https://docs.lytics.com/docs/account-users https://docs.lytics.com/docs/privacy-and-data-protection https://docs.lytics.com/discuss/673d9424e0ae0a001251fe54
12. Does your solution offer built-in consent & preference management that suppresses activation instantly when permissions change (GDPR/CCPA ready)?	Yes	Yes. The platform provides a robust framework for managing consent and preferences by ingesting user consent signals from external platforms or via its own APIs, but it functions as a consent enforcement and propagation engine rather than a self-contained, built-in consent management platform. The platform is designed to integrate with third-party Consent Management Platforms (CMPs) like OneTrust. These consent attributes are mapped to the unified profile and used to enforce data suppression across activation channels in near real-time, supporting compliance with regulations like GDPR and CCPA.	https://docs.lytics.com/docs/managing-consumer-data-subject-requests-with-onetrust https://docs.lytics.com/docs/privacy-and-data-protection

Oct 2025 CDP Pulse Report - Contentstack

Yes/No	Details	Source(s)
13. Does your solution support regional data residency options (e.g., EU-only, US-only) with encryption at rest (AES-256) and in transit (TLS 1.2+)?	Yes. Contentstack CDP (Lytics) supports regional data residency, with options to host accounts in the United States or the European Union and the data is isolated per region (no cross-region mixing). The platform encrypts data at rest and in transit using HTTPS/TLS for data in transit and AES-256 encryption for data at rest (and backups), ensuring compliance with security best practices.	https://www.contentstack.com/trust https://docs.lytics.com/discuss/673d9424e0ae0a001251fe54 https://docs.lytics.com/discuss/64540892cd9e0022f66c0c0
14. Does your solution provide reverse-ETL / audience-syndication connectors that push profiles and events to major ad, analytics, and messaging destinations natively?	Yes. Contentstack CDP (Lytics) provides robust reverse-ETL and audience syndication capabilities through its library of pre-built integrations. These "Jobs" are configured to push user profiles, computed attributes, and audience memberships to a wide array of third-party advertising, analytics, and messaging destinations.	https://docs.lytics.com/docs/data-sources https://docs.lytics.com/docs/integrated-marketing-tools https://docs.lytics.com/docs/lytics-integration-options
15. Does your solution expose a unified customer-profile API/stream that any downstream system can query or subscribe to in real time?	Yes. The platform exposes a real-time, RESTful Profile API that allows any authenticated downstream system to query a complete, unified profile on demand. Furthermore, it can stream profile updated events and audience membership changes to external systems via integrations with message queues like Amazon Kinesis for event-driven architectures. For example, it supports exporting an individual's profile as a JSON object through its API. In addition, webhook and streaming integrations are available for pushing real-time profile/event updates to subscribers, enabling other services to consume up-to-date customer data.	https://docs.lytics.com/docs/privacy-and-data-protection https://docs.lytics.com/docs/data-streams-1 https://docs.lytics.com/docs/aws-kinesis-overview
16. Does your solution publish transparent, usage-based pricing (events, MTU, storage) in writing, independent of professional-services spend?	Yes. Lytics (Contentstack's CDP) publicly lists a usage-based pricing model. The pricing is credit-based (tied to data volumes like events/profiles) with clearly published tiers. For example, Lytics offers a free tier and a "Growth" plan (~\$500/month for 5 million monthly credits), and these rates are documented on its website. This pricing disclosure is separate from any services and is transparent to customers.	https://www.lytics.com/pricing/

Vendor Deep Dive



Session AI

Primary Archetype: Independent CDP (Marketer-led)

Strengths: Warehouse-native/Composable

Capability Lens: Data Visibility & Quality; Security, Consent & Residency; Customization & Composability

VOX Score: 69.14 / 100

Preference Coverage: 11 / 16 questions “Yes”

📞 Talk to Session AI: <https://www.sessionai.com/>

📞 [Ask VOX for Introduction](#)

Oct 2025 CDP Pulse Report - Session AI

Yes/No	Details	Source(s)
1. Does your solution provide data-observability dashboards that refresh in ≤ 60 seconds (p95) and show source-to-destination lineage for every first-party data stream?	No	Session AI offers real-time insight dashboards (with auto-refresh options in seconds) for monitoring campaign performance, but it does not highlight dedicated data observability or end-to-end lineage tracing for each data stream. The platform's focus is on in-session metrics and actions rather than full data pipeline lineage.
2. Does your solution ship native inbound connectors or APIs for the top-10 commerce, marketing-automation, and service platforms without custom code or PS?	Yes	Session AI provides pre-built connectors and flexible APIs to ingest data from popular commerce and marketing systems without heavy custom code. For example, it offers native integrations for platforms like Salesforce Data 360, Adobe Commerce (Magento), HCL Commerce, and Adobe Campaign, and supports data ingestion via webhooks, JavaScript SDK, REST API, and batch uploads. These out-of-the-box connectors minimize the need for custom development or professional services.
3. Does your solution support server-side and client-side collection for web, mobile-app, email, and in-store (POS/scanner) events in a single SDK?	Yes	The Session AI platform supports both client-side and server-side data collection across channels. It offers SDKs for web (JavaScript), iOS, Android, Windows, and more, alongside server-side endpoints (webhooks, REST APIs, batch collectors) to capture events. This unified approach allows collection of website events, mobile app interactions, and even offline or back-end events (like POS transactions or email responses) into the same system, ensuring all first-party event data flows into one platform.
4. Does your solution offer deterministic identity-resolution that stitches online cookies, MAIDs, loyalty IDs and POS receipts into one customer profile?	No	Session AI does not focus on traditional identity resolution across channels. In fact, it emphasizes the ability to personalize and predict behavior without needing to stitch personal identifiers. The platform works in a privacy-first manner, analyzing in-session behavioral signals without relying on PII or merging profiles from cookies, device IDs, loyalty accounts, etc. In practice, it avoids identity stitching and instead segments and targets anonymous visitors based on real-time behavior alone.
5. Does your solution sync online and in-store inventory events in < 5 minutes (p95) to enable BOPIS and endless-aisle use cases?	Yes	Session AI is designed for real-time event processing and can incorporate inventory data quickly to support omnichannel use cases. The platform notes that as brands mature with in-session marketing, they use it to orchestrate actions like inventory clearance and BOPIS (Buy Online, Pick Up In Store). This implies the system can ingest updates such as inventory changes from in-store systems and reflect them or act on them in near real-time (on the order of minutes or faster), enabling endless-aisle and BOPIS scenarios where up-to-date inventory information is crucial.

Oct 2025 CDP Pulse Report - Session AI

Yes/No	Details	Source(s)	
6. Does your solution expose decision or personalization APIs that return in < 1 second (p95) so actions fire while the session is active?	Yes	Session AI's decisioning engine operates extremely quickly. It uses machine learning to score a visitor's purchase intent within their first few clicks and can deliver a personalized action in real-time. In fact, the end-to-end process from observing behavior to selecting and returning an in-session action is stated to take under 100 milliseconds, well below a second. This sub-second latency ensures that personalized offers or recommendations are returned almost instantly, allowing them to be presented while the user is still active on the site.	https://www.sessionai.com/capabilities/platform-overview/
7. Does your solution include embedded predictive-merchandising models (e.g., next-best product) deployable by business users without data-science help?	Yes	Session AI comes with built-in AI models that business users can leverage without coding or data science expertise. Notably, it has a patented "purchase propensity" model that predicts a visitor's likelihood to buy in real time, as well as templates for use cases like on-the-fence shoppers and product urgency. These out-of-the-box predictive models (e.g., to identify who needs an incentive or which products are trending) are accessible through a visual Experience Gallery, allowing marketers to deploy machine learning-driven experiences (such as targeted offers or recommendations) without needing to build models from scratch.	https://www.sessionai.com/capabilities/platform-overview/ https://www.sessionai.com/blog/empowering-business-users-to-build-powerful-customer-experiences/
8. Does your solution allow teams to upload and deploy their own ML/AI models (PMML/ONNX/Python) to score profiles or events in real time?	Yes	The Session AI platform provides a framework for custom machine learning models and real-time scoring pipelines. It includes a "Predictions" module where users can create and deploy custom models, define feature libraries, and set up model pipelines and validators. This indicates that teams can bring their own algorithms or tailor models to their needs within the platform. While specific formats (PMML/ONNX) aren't explicitly cited, the ability to define custom model code and integrate it into Session AI's real-time decision flow suggests that organizations can deploy their own ML logic to score events or profiles on the fly.	https://docs.zineone.com/docs/working-with-machine-learning https://docs.zineone.com/docs/working-with-machine-learning
9. Does your solution provide an SQL workspace or direct lake-house access for segmentation, attribution, and ad-hoc analytics?	Yes	Session AI offers tools for advanced analytics on collected data. Users have access to an ad-hoc query interface (within the platform's Insights section) where they can query event and profile data and even build custom charts and dashboards. The platform also supports creation of data cubes and running analytical queries on event streams. While it's an in-platform workspace (not a direct connection to an external lakehouse), it enables analysts to perform segmentation, attribution analysis, and other custom queries using a SQL-like approach on the first-party data captured by Session AI.	https://docs.zineone.com/docs/creating-an-ad-hoc-query https://docs.zineone.com/docs/creating-dashboard-designs

Oct 2025 CDP Pulse Report - Session AI

Yes/No	Details	Source(s)
No	<p>10. Does your solution run continuous data-quality and anomaly detection with automatic quarantine or alerting before bad data feeds downstream systems?</p> <p>There is no indication that Session AI includes an automated data quality monitoring or anomaly detection system that quarantines bad data. The platform documentation does not describe features for ongoing anomaly detection on incoming events or automatic blocking of suspicious data. Data validations in Session AI appear to be something that can be configured manually (for instance, developers can set up custom "validators" or pre-processors in the data pipeline), but there isn't a built-in continuous anomaly detection service actively preventing bad data from propagating.</p>	https://docs.zineone.com/docs/creating Validators
Yes	<p>11. Does your solution enforce role-based access controls and immutable audit logs at the field level across all regions?</p> <p>Session AI provides enterprise-grade administrative controls including role-based user management and auditing. Administrators can create and manage users for the Session AI portal (with the ability to configure SSO and password policies), implying definable roles or permissions. The platform also maintains audit logs – for example, it offers a "Master Audit Trail" that records changes or events in the system. While field-level permission granularity isn't explicitly detailed, the presence of these security and admin features and the fact that the platform can be deployed for different regions suggest that access can be controlled per role and all critical changes are logged in an immutable audit trail for compliance.</p>	https://docs.zineone.com/docs/creating Dashboard Designs
Yes	<p>12. Does your solution offer built-in consent & preference management that suppresses activation instantly when permissions change (GDPR/CCPA ready)?</p> <p>Session AI is built with privacy compliance in mind and includes features to honor user consent choices. It provides APIs for consent management – for instance, there are endpoints to handle consumer privacy requests like profile opt-outs (e.g., to comply with CCPA "Do Not Sell" or GDPR withdrawal of consent). Using these, a user's data can be flagged such that the system will immediately stop using that profile for personalization/activation once an opt-out is received. In practice, this means that if a customer revokes consent, Session AI can instantly suppress any further marketing actions for that individual, supporting GDPR/CCPA requirements.</p>	https://devguide.zineone.com/docs/introduction-to-developers-guide
Yes	<p>13. Does your solution support regional data residency options (e.g., EU-only, US-only) with encryption at rest (AES-256) and in transit (TLS 1.2+)?</p> <p>Session AI's infrastructure is deployed on secure cloud services (like AWS and Azure), which means it can be hosted in specific regions to meet data residency needs (for example, keeping European customer data in EU data centers). The platform adheres to industry-standard security practices: data is transmitted over TLS and stored with strong encryption. (While specific encryption protocols aren't listed on the marketing site, AWS infrastructure generally ensures AES-256 encryption at rest and TLS 1.2+ for data in transit, and Session AI's availability on these clouds indicates compliance with those security standards.)</p>	https://www.sessionai.com/integrations-connectors/

Oct 2025 CDP Pulse Report - Session AI

Yes/No	Details	Source(s)	
14. Does your solution provide reverse-ETL / audience-syndication connectors that push profiles and events to major ad, analytics, and messaging destinations natively?	No	Session AI is primarily an in-session personalization and decisioning platform and not a full customer data platform for broad audience syndication. It does have a few integrations to export data (for example, it can trigger campaigns in Adobe Campaign or share intelligence with Pega systems), and data can always be pulled via its APIs. However, it doesn't offer a wide array of native "reverse ETL" connectors to automatically sync audiences into the likes of Facebook Ads, Google Ads, or various CRM and analytics tools out-of-the-box. Any such exports would require using the available APIs or custom solutions, as the platform's native connectors focus more on ingesting data and enabling real-time on-site use cases rather than batch syndication to external marketing channels.	https://www.sessionai.com/integrations-connectors/ https://www.sessionai.com/integrations-connectors/
15. Does your solution expose a unified customer-profile API/stream that any downstream system can query or subscribe to in real time?	Yes	Session AI maintains unified profiles for visitors (capturing their behaviors, context, and scores) and provides programmatic access to this data. It offers RESTful APIs through which external systems can query profile details or even fetch live customer activity streams. For example, there are endpoints to retrieve a customer's profile attributes and recent events. This means downstream systems can pull Session AI's intelligence (such as propensity scores or segment assignments) on demand, enabling real-time integration of Session AI's profile data into other tools or databases.	https://devguide.zineone.com/docs/introduction-to-developers-guide
16. Does your solution publish transparent, usage-based pricing (events, MTU, storage) in writing, independent of professional-services spend?	No	Session AI's pricing is not openly published as a usage-based model. The pricing is typically obtained via sales consultation and is structured as enterprise contracts rather than self-service metered rates. There is no public price list for events, MTUs, or storage consumption. In fact, indications from partner marketplaces suggest a large annual license cost – for example, a Salesforce AppExchange listing notes a starting price around \$250k per year for a site's traffic – which underscores that pricing is negotiated and not transparently usage-tied or posted online.	https://appexchange.salesforce.com/appxListingDetail?listingId=5eae98a9-587a-4941-aed3-5801ade884b2 https://aws.amazon.com/marketplace/pp/prodview-k5lvtcuifgpk

Vendor Deep Dive



Amperity

Primary Archetype: Independent CDP (Marketer-led)

Strengths: Retail Specialist; Governance

Capability Lens: Data Visibility & Quality; Security, Consent & Residency; Customization & Composability

VOX Score: 87.22 / 100

Preference Coverage: 14 / 16 questions “Yes”

📞 Talk to Amperity: <https://amperity.com/>

📞 [Ask VOX for Introduction](#)

Oct 2025 CDP Pulse Report - Amperity

Yes/No	Details	Source(s)
1. Does your solution provide data-observability dashboards that refresh in ≤ 60 seconds (p95) and show source-to-destination lineage for every first-party data stream?	Yes Amperity offers data observability features including data lineage graphs, workflow monitoring, and usage dashboards, all of which refresh and update in real-time each time the user loads the page. While Amperity's documentation doesn't explicitly state a formal SLA (e.g., a ≤60-second refresh guarantee at the 95th percentile), according to their team, these observability views consistently refresh upon page load, effectively providing immediate updates.	https://docs.amperity.com/reference/data_lineage.html
2. Does your solution ship native inbound connectors or APIs for the top-10 commerce, marketing-automation, and service platforms without custom code or PS?	Yes Amperity provides hundreds of pre-built native connectors for popular commerce, marketing, and service platforms. For example, it has built-in integrations for e-commerce and marketing systems like Adobe/Magento Commerce, Shopify, Braze, Cordial, HubSpot, Mailchimp, Salesforce Marketing Cloud, Zendesk, etc., which can be used out-of-the-box without custom code. These connectors enable data ingestion from and activation to top tools without requiring professional services.	https://docs.amperity.com/reference/connected_systems.html
3. Does your solution support server-side and client-side collection for web, mobile-app, email, and in-store (POS/scanner) events in a single SDK?	Yes Amperity supports both client-side and server-side event collection across channels, though via multiple tools rather than a single monolithic SDK. It offers a client-side JavaScript tag/SDK for web data capture (configurable with an Amperity write token and stream ID), and a Streaming Ingest REST API for server-side or mobile app event streaming in real time. Through these APIs/SDKs, brands can ingest web behaviors, mobile app events, email engagement, and offline/POS transactions into Amperity's platform. (External analyses note that Amperity can collect events via SDK and resolve identities across client and server inputs.)	https://docs.amperity.com/reference/api.html
4. Does your solution offer deterministic identity-resolution that stitches online cookies, MAIDs, loyalty IDs and POS receipts into one customer profile?	Yes Identity resolution is a core feature of Amperity. Its "Stitch" engine uses AI/ML (probabilistic and deterministic methods) to unify customer records across all identifiers and touchpoints into a single profile. This includes linking online identifiers (e.g. first-party web cookies, mobile device IDs) with offline identifiers (loyalty program IDs, email addresses, point-of-sale transaction IDs/receipts, etc.) into one comprehensive profile. Amperity's patented identity resolution merges all customer interactions (online and in-store) into unified profiles, although it emphasizes machine-learning-based matching rather than solely rule-based deterministic keys.	https://www.businesswire.com/news/home/20230518005375/en/Amperity-Replaces-Third-Party-Cookies-to-Drive-Customer-Acquisition-and-ROI
5. Does your solution sync online and in-store inventory events in < 5 minutes (p95) to enable BOPIS and endless-aisle use cases?	Yes Amperity can handle near-real-time data ingestion and availability. It supports streaming data feeds such as inventory updates with low latency – for example, data streamed via the Streaming Ingest API appears in Amperity's real-time tables typically within about 2 minutes. This sub-5-minute data propagation allows use cases like BOPIS (Buy Online Pickup In Store) and "endless aisle," as inventory events from online or in-store systems can be reflected in the customer 360 and activation workflows almost immediately. (Real-time tables complement the daily batch processing to ensure fresh event data is usable within minutes.)	https://docs.amperity.com/operator/realtime.html
6. Does your solution expose decision or personalization APIs that return in < 1 second (p95) so actions fire while the session is active?	Yes Amperity provides a low-latency Profile API for personalization and decisions. The Profile API endpoints are backed by indexed customer profiles and are designed for fast lookup performance. In benchmark tests, even fairly large profile indexes (millions of rows) showed 90th-percentile response times around ~100–200ms and 99th-percentile well below 1 second. This means personalized data (e.g. customer attributes or product recommendations) can be retrieved via API in a fraction of a second, enabling in-session actions on websites or apps while the user is still active.	https://docs.amperity.com/operator/api_profile.html
7. Does your solution include embedded predictive-merchandising models (e.g., next-best product) deployable by business users without data-science help?	Yes Amperity's platform comes with built-in predictive models (branded as AmpAI) that non-technical users can leverage. These include models like product affinity (to predict next best product or category for each customer), churn propensity, customer lifetime value, and others. Business users can deploy these out-of-the-box machine learning models through the UI – for example, generating scores or segments for "next-best product" recommendations – without needing to build models from scratch. The system provides these predictive insights (e.g., product recommendations, likelihood to churn) natively, which can be used in campaigns or passed to other systems.	https://amperity.com/platform/customer-360

Oct 2025 CDP Pulse Report - Amperity

	Yes/No	Details	Source(s)
8. Does your solution allow teams to upload and deploy their own ML/AI models (PMML/ONNX/Python) to score profiles or events in real time?	No	Amperity does not support importing or deploying custom-trained ML models (e.g., PMML, ONNX, or Python artifacts) directly within its platform for real-time scoring. Instead, users typically leverage their own data warehouses and modeling tools to build and score proprietary models externally. Through Amperity Bridge—a live connection to the customer's data warehouse—users can instantly visualize and utilize the outputs from these externally-run models within Amperity, eliminating the need for manual result ingestion.	https://docs.amperity.com/reference/bridge.html
9. Does your solution provide an SQL workspace or direct lake-house access for segmentation, attribution, and ad-hoc analytics?	Yes	Amperity offers both an internal SQL workspace and the ability to connect with external "lakehouse" analytics environments. The platform includes a SQL query editor (using Presto SQL) where analysts can write queries against the unified customer 360 database or any published tables for ad-hoc analysis and segmentation. Moreover, Amperity has a feature called Amperity Bridge that enables sharing data with cloud data warehouses/data lakes (e.g., Snowflake or Databricks) without data replication. This means companies can directly access and query Amperity-managed customer data in their own lakehouse environment for advanced analytics or attribution modeling.	https://docs.amperity.com/amp360/workflows.html https://docs.amperity.com/reference/bridge.html
10. Does your solution run continuous data-quality and anomaly detection with automatic quarantine or alerting before bad data feeds downstream systems?	Yes	Amperity includes robust data quality monitoring and alerts to catch bad data before it causes issues downstream. It provides Input Validation Reports that check the integrity of incoming data (e.g., the percentage of missing or malformed emails, phone numbers, transaction fields). These reports help identify anomalies or data quality problems early, and the platform surfaces alerts/notifications in the UI if workflows encounter errors or unusual patterns. While Amperity will inform users of data anomalies (and allows configuring rules to ignore or flag "bad values"), it does not outright halt all pipelines automatically; instead it quarantines issues by alert and gives the team an opportunity to fix data before activation. (No evidence of fully automatic quarantine, but proactive alerts and QA checks are in place.)	https://docs.amperity.com/operator/validate_inputs.html https://docs.amperity.com/amp360/workflows.html
11. Does your solution enforce role-based access controls and immutable audit logs at the field level across all regions?	Yes	Amperity provides enterprise-grade security with granular role-based access control (RBAC) and comprehensive auditing. User access can be restricted such that users only see the fields/data appropriate for their role, with personally identifiable info (PII) masked or redacted based on permissions. All user actions and changes in the system are transparently logged – Amperity maintains audit logs of activities (ensuring "transparent access to user actions" for compliance). These audit trails are immutable records. The platform is SOC 2 certified and supports multi-region deployments, so security and access controls apply uniformly across regions (with options for regional data hosting as needed).	https://amperity.com/platform/reverse-etl https://amperity.com/platform/reverse-etl
12. Does your solution offer built-in consent & preference management that suppresses activation instantly when permissions change (GDPR/CCPA ready)?	Yes	Amperity includes built-in consent and privacy management features to honor customer permissions in real-time. It has automated workflows for privacy requests (such as GDPR/CCPA Data Subject Access/Deletion Requests) and can update or suppress data based on consent flags. For example, if a customer revokes consent or opts out, Amperity's consent management will ensure that the customer's data is excluded from any new audience activation immediately. The platform provides auditable consent management and compliance tools, and it integrates preference/opt-out statuses into segmentation logic so that campaigns automatically suppress any profiles that should not be contacted. (Amperity's documentation on Privacy Rights confirms it is GDPR/CCPA-ready with built-in consent handling.)	https://amperity.com/platform/reverse-etl
13. Does your solution support regional data residency options (e.g., EU-only, US-only) with encryption at rest (AES-256) and in transit (TLS 1.2+)?	Yes	Amperity can be deployed in various regions to meet data residency requirements – it offers multiple cloud hosting options so that data can reside in specific jurisdictions (e.g., an EU-only instance for European data). All customer data is secured with strong encryption: data is encrypted at rest using 256-bit AES and in transit via modern TLS (Amperity notes encryption upon data ingest, at rest, and upon egress). The platform meets enterprise security standards (SOC 2, HIPAA, GDPR compliance), indicating that encryption and regional isolation of data are part of its capabilities.	https://amperity.com/platform/reverse-etl https://amperity.com/platform/reverse-etl
14. Does your solution provide reverse-ETL / audience-syndication connectors that push profiles and events to major ad, analytics, and messaging destinations natively?	Yes	Amperity has native audience syndication (activation) connectors to a wide range of marketing, advertising, and analytics platforms. Users can easily push customer profiles or event segments from Amperity to dozens of downstream systems without custom development. For example, Amperity's UI supports sending audiences to advertising platforms (Google Ads, Facebook/Meta, TikTok, Snapchat, The Trade Desk, DV360, etc.), email/SMS and marketing automation tools (Salesforce Marketing Cloud, Adobe Campaign, Braze, Klaviyo, Mailchimp, etc.), analytics or data destinations (Snowflake, S3, Google BigQuery), and many more. These connectors are maintained by Amperity and allow "point-and-click" export of segments and event data to all major channels.	https://docs.amperity.com/ampig/models.html https://docs.amperity.com/ampig/models.html

Oct 2025 CDP Pulse Report - Amperity

Yes/No	Details	Source(s)	
15. Does your solution expose a unified customer-profile API/stream that any downstream system can query or subscribe to in real time?	Yes	<p>Amperity provides a unified profile API that external systems can query in real-time for customer data. The Profile API allows the creation of custom REST endpoints on top of the unified customer 360 database. Downstream applications (e.g., a website, mobile app, or service platform) can call these endpoints with a customer identifier to retrieve that customer's latest profile attributes (or even subscribe to changes via polling or event integration). This enables real-time access to the single customer view from any channel. Additionally, Amperity's architecture supports streaming updates (via real-time tables and API) so that subscribers always get up-to-date profile information during active sessions.</p>	https://docs.amperity.com/operator/api_profile.html
16. Does your solution publish transparent, usage-based pricing (events, MTU, storage) in writing, independent of professional-services spend?	No	<p>Amperity's pricing model is usage-based, measured in units called 'Amps,' though specific pricing metrics or rate cards are not publicly disclosed. Pricing typically aligns with monthly tracked users (MTUs) and data volumes. All subscriptions include implementation onboarding, training materials, sandboxes, and SLA-backed support. Customers receive a dedicated implementation team as part of their subscription, with optional professional services available as add-ons for deeper ongoing support.</p>	https://amperity.com/pricing

Vendor Deep Dive



BlueConic

Primary Archetype: Independent CDP (Marketer-led)

Strengths: Real-time/Edge

Capability Lens: Data Visibility & Quality; Security, Consent & Residency; Customization & Composability

VOX Score: 100 / 100

Preference Coverage: 16 / 16 questions “Yes”

📞 Talk to BlueConic: <https://www.blueconic.com/>

📞 [Ask VOX for Introduction](#)

Oct 2025 CDP Pulse Report - BlueConic

Yes/No	Details	Source(s)
1. Does your solution provide data-observability dashboards that refresh in ≤ 60 seconds (p95) and show source-to-destination lineage for every first-party data stream?	Yes BlueConic provides robust, customizable reporting dashboards and insights that offer real-time visibility into data flow, including source-to-destination monitoring. While BlueConic does not explicitly document a guaranteed sub-60-second SLA, many dashboard insights refresh automatically upon page load—some updating every second—supporting near-real-time (<60s) monitoring in practice. Additionally, BlueConic's dashboards include high-level data-flow diagrams, connection overviews, batch-run histories, and granular profile-origin tracking, enabling users to visualize and trace data lineage effectively.	https://support.blueconic.com/en/collectio...ns/ https://support.blueconic.com/en/articles/ https://support.blueconic.com/en/articles/ https://support.blueconic.com/en/articles/ https://support.blueconic.com/en/
2. Does your solution ship native inbound connectors or APIs for the top-10 commerce, marketing-automation, and service platforms without custom code or PS?	Yes BlueConic offers over 100 pre-built, native connectors to major commerce, marketing automation, CRM, advertising, analytics, and data platforms—such as Shopify, Magento, Salesforce, Marketo (Adobe), Eloqua, Google Analytics, Amazon Ads, and many others. Each connector provides configurable, standardized profile property mappings directly within the user interface, enabling non-technical users to easily integrate systems without custom coding or professional services. Additionally, BlueConic includes built-in data processing tools (e.g., normalization, validation, cleansing) to ensure imported data is accurate and actionable, all configurable through intuitive UI settings.	www.blueconic.com/connections https://support.blueconic.com/en/articles/247657-data-processors-overview
3. Does your solution support server-side and client-side collection for web, mobile-app, email, and in-store (POS/scanner) events in a single SDK?	Yes BlueConic provides comprehensive multi-channel data collection capabilities, enabling unified customer profiles across web, mobile apps, email, and offline (in-store/POS) interactions. It achieves this through a combination of client-side SDKs (JavaScript snippet for web, native SDKs for iOS, Android, React Native apps), server-side REST APIs, and extensive configurable connectors. These methods support real-time and batch data ingestion across all channels, effectively unifying behavioral, demographic, transactional, and consent data into a single customer profile. Additionally, built-in data listeners provide real-time enrichment directly on digital channels, while server-to-server connections and APIs seamlessly integrate offline events such as POS transactions into BlueConic's unified customer view.	https://support.blueconic.com/en/ https://www.blueconic.com/connections https://rest.apidoc.blueconic.com/ https://support.blueconic.com/en/articles/248051-javascript-front-end-api
4. Does your solution offer deterministic identity-resolution that stitches online cookies, MAIDs, loyalty IDs and POS receipts into one customer profile?	Yes BlueConic provides robust identity resolution, merging customer profiles across multiple channels (web, mobile, email, offline) through both deterministic and probabilistic matching methods. Deterministic merges use explicit identifiers—such as email, mobile Ad IDs, loyalty numbers, POS transaction IDs, or any user-defined unique attribute—to immediately unify profiles when a match is identified. Additionally, BlueConic supports advanced probabilistic ('fuzzy') matching via its AI Workbench, enabling customizable scoring based on weighted profile attributes (e.g., names, phone numbers), automatically merging profiles that meet client-defined confidence thresholds. Once merged, the unified profiles persist server-side, maintaining continuous identity linkage across channels and devices.	https://support.blueconic.com/en/articles/ https://support.blueconic.com/en/articles/ https://www.blueconic.com/resources/

Oct 2025 CDP Pulse Report - BlueConic

Yes/No	Details	Source(s)	
5. Does your solution sync online and in-store inventory events in < 5 minutes (p95) to enable BOPIS and endless-aisle use cases?	Yes	BlueConic supports real-time and near-real-time synchronization of customer data, making it ideal for use cases like BOPIS (Buy Online, Pick Up In-Store) and endless aisle scenarios. The platform provides multiple real-time data ingestion methods—including SDKs, REST APIs, and configurable triggers (webhooks, listeners)—to ingest and propagate inventory updates from external systems (such as POS, OMS, or inventory management) almost instantly. Additionally, BlueConic's flexible batch connections allow frequent micro-batching intervals (down to every few minutes), further enhancing data freshness. Although BlueConic does not publish a formal “<5-minute” SLA, it is explicitly designed to propagate inventory and catalog updates within minutes, keeping online and in-store experiences consistently synchronized.	https://www.blueconic.com/connections
6. Does your solution expose decision or personalization APIs that return in < 1 second (p95) so actions fire while the session is active?	Yes	BlueConic is designed explicitly for in-session, real-time personalization and decisioning. Its embedded JavaScript snippet (just <40KB, cached after the initial load via CloudFront CDN) powers immediate, sub-second content personalization during active user sessions. Additionally, BlueConic provides a comprehensive REST API suite for real-time profile and segment queries, with internal benchmarks showing typical API responses within 500ms under optimal conditions—though actual latency may vary due to external network factors. Developers have access to extensive APIs and SDKs (JavaScript, Python, iOS, Android, React Native) enabling flexible, real-time interactions and personalization.	https://rest.apidoc.blueconic.com/ https://support.blueconic.com/en/collections/609503-developer-tools
7. Does your solution include embedded predictive-merchandising models (e.g., next-best product) deployable by business users without data-science help?	Yes	BlueConic offers extensive AI and predictive modeling capabilities out of the box, including pre-built models for customer lifetime value (CLV), RFM scoring, churn propensity, next-best actions (NBA), and product recommendations—all easily configurable by non-technical users through intuitive, no-code interfaces. Additionally, BlueConic's AI Workbench provides advanced users and data scientists with integrated Jupyter notebooks (supporting Python) to customize existing models or develop entirely new predictive models. This hybrid approach combines powerful built-in models accessible via the UI with a flexible, extensible Python environment, allowing users to enhance predictive insights and personalize customer experiences.	https://support.blueconic.com/en/articles/247838-ai-workbench-overview
8. Does your solution allow teams to upload and deploy their own ML/AI models (PMML/ONNX/Python) to score profiles or events in real time?	Yes	BlueConic's AI Workbench supports custom machine-learning models through an integrated Jupyter notebook environment, allowing data scientists to directly write, train, and execute Python code within the platform. Teams can import their own Python notebooks or ML models (including PMML or ONNX formats via Python libraries) to enrich profiles and run scoring jobs at flexible intervals—from real-time updates to frequent batch schedules. Model outputs immediately enrich BlueConic's unified customer profiles, supporting segmentation, personalization, and analytics use cases. This hybrid capability uniquely combines robust out-of-the-box models accessible via a no-code UI with customizable, user-defined models for advanced predictive analytics and real-time activation.	https://support.blueconic.com/en/articles/ https://supportdocs.blueconic.com/python/

Oct 2025 CDP Pulse Report - BlueConic

Yes/No	Details	Source(s)
9. Does your solution provide an SQL workspace or direct lake-house access for segmentation, attribution, and ad-hoc analytics?	Yes BlueConic provides SQL-based segmentation, attribution, and ad-hoc analytics by seamlessly exporting data to external lake-house environments (e.g., Snowflake, BigQuery). Within the platform, users can leverage comprehensive REST APIs, SDKs, and the integrated AI Workbench (Python/Jupyter notebooks) for advanced analytics.	https://rest.apidoc.blueconic.com/
10. Does your solution run continuous data-quality and anomaly detection with automatic quarantine or alerting before bad data feeds downstream systems?	Yes BlueConic includes features to ensure data quality, though they may require some configuration. The platform has built-in data cleansing and normalization (including ML techniques to maintain data hygiene) so that profile data stays accurate. Additionally, BlueConic provides a specialized Data Validation connector that can be used to detect anomalies or mismatches in incoming data before it's applied to profiles. For example, you can use this to compare an imported dataset vs. what's in BlueConic and catch any discrepancies (like missing or malformed records) and then prevent those bad records from contaminating downstream profile data. While not an autonomous "self-healing quarantine" across every stream, these tools allow teams to implement automated data checks and receive alerts or hold back bad data if validation rules fail, thereby protecting downstream systems from faulty data.	https://support.blueconic.com/en/collections/ www.blueconic.com/connections https://support.blueconic.com/en/articles/ https://support.blueconic.com/
11. Does your solution enforce role-based access controls and immutable audit logs at the field level across all regions?	Yes BlueConic supports robust security controls including role-based access control (RBAC) down to data field sensitivity. Administrators can define user roles and permissions such that certain roles cannot view or export PII or specific profile properties. This effectively restricts access to fields (like names, emails, etc.) based on role. BlueConic also maintains audit logs for system and user activities. It provides an Audit Event API that allows extraction of a log of events for security monitoring. All changes or accesses in the platform can thus be recorded and monitored. With regional hosting, these controls and logs apply across all data regions, ensuring compliance and traceability globally. (The audit logs are immutable in the sense that every event is recorded and can be fed to a SIEM for an indelible record of who did what.)	https://support.blueconic.com/en/articles/ https://rest.apidoc.blueconic.com/
12. Does your solution offer built-in consent & preference management that suppresses activation instantly when permissions change (GDPR/CCPA ready)?	Yes BlueConic has built-in consent and preference management capabilities to comply with privacy laws like GDPR and CCPA. It allows you to define consent categories ("objectives") and will automatically enforce them in real time. For instance, if a user withdraws consent, BlueConic immediately suppresses any data collection or activation related to that consent category. Consent statuses are stored per profile, and all BlueConic's connectors and listeners check those before executing. The platform supports multi-region privacy rules (opt-in vs opt-out regimes) and will instantly update a visitor's eligibility for campaigns when their preferences change. In short, when a user's consent preference is updated, BlueConic propagates that change instantly across the system to prevent any unauthorized data use or activation.	https://support.blueconic.com/hc/en-us/ https://support.blueconic.com/hc/en-us/ https://support.blueconic.com/hc/en-us/

Oct 2025 CDP Pulse Report - BlueConic

	Yes/No	Details	Source(s)
13. Does your solution support regional data residency options (e.g., EU-only, US-only) with encryption at rest (AES-256) and in transit (TLS 1.2+)?	Yes	BlueConic offers regional hosting and strong encryption. Customers can choose the region where their BlueConic instance/data is hosted (for example, EU data can be kept on EU servers, US data in US, etc.), and BlueConic ensures data stays in that region. In terms of security, all data is encrypted at rest using AES-256 encryption. In transit, BlueConic uses HTTPS with modern TLS (versions 1.2/1.3) and at least 128-bit encryption keys for all communications. This means both regional data residency compliance and enterprise-grade encryption standards are in place by default (along with SOC 2 certified security practices).	https://support.blueconic.com/en/articles/
14. Does your solution provide reverse-ETL / audience-syndication connectors that push profiles and events to major ad, analytics, and messaging destinations natively?	Yes	BlueConic has an extensive library of outbound connectors for audience syndication (essentially reverse-ETL). It can natively push customer profiles, segments, or event data to all major marketing and analytics platforms. For example, BlueConic includes connectors for advertising networks (Facebook Ads, Google Ads, LinkedIn, TikTok, etc.), analytics tools (Google Analytics 4, Adobe Analytics, Snowflake, BigQuery), and messaging/ESP systems (such as Mailchimp, Salesforce Marketing Cloud, Klaviyo, HubSpot, etc.). These integrations are built-in and require no custom development – you simply configure credentials and mapping, and BlueConic will continuously export or sync audience data to the chosen downstream platforms.	http://www.blueconic.com/connections
15. Does your solution expose a unified customer-profile API/stream that any downstream system can query or subscribe to in real time?	Yes	BlueConic provides both APIs and real-time streaming for unified profiles. There is a RESTful BlueConic Customer Profile API that allows external systems to query profile data, segments, and even listen for interaction events. In addition, BlueConic offers “firehose” streaming connections which can send profile updates and events in real time to message queuesstreams (like AWS Kinesis, Google Pub/Sub, Azure Event Hubs). This means downstream systems can either call an API to fetch the latest profile attributes on demand, or subscribe to a live feed of profile-event data. Both methods provide access to the unified customer profile data in real time for any application that needs it.	https://rest.apidoc.blueconic.com/ https://support.blueconic.com/en/collections/609503-developer-tools
16. Does your solution publish transparent, usage-based pricing (events, MTU, storage) in writing, independent of professional-services spend?	Yes	BlueConic provides both APIs and real-time streaming for unified profiles. There is a RESTful BlueConic Customer Profile API that allows external systems to query profile data, segments, and even listen for interaction events. In addition, BlueConic offers “firehose” streaming connections which can send profile updates and events in real time to message queuesstreams (like AWS Kinesis, Google Pub/Sub, Azure Event Hubs). This means downstream systems can either call an API to fetch the latest profile attributes on demand, or subscribe to a live feed of profile-event data. Both methods provide access to the unified customer profile data in real time for any application that needs it.	https://rest.apidoc.blueconic.com/ https://support.blueconic.com/en/collections/609503-developer-tools

Commerce & Personalization Engines



Vendor Deep Dive



Iterable

Primary Archetype: Commerce & Personalization Engine

Strengths: Engagement Hub; Retail Specialist;
Real-time/Edge

Capability Lens: In-Session Activation

VOX Score: 56.16 / 100

Preference Coverage: 9 / 16 questions “Yes”

📞 Talk to Iterable: <https://iterable.com/>

📞 [Ask VOX for Introduction](#)

Oct 2025 CDP Pulse Report - Iterable

Yes/No	Details	Source(s)
No	<p>1. Does your solution provide data-observability dashboards that refresh in ≤ 60 seconds (p95) and show source-to-destination lineage for every first-party data stream?</p> <p>Iterable offers campaign performance dashboards and monitors data ingestion jobs, but it does not include a dedicated real-time data observability dashboard with sub-minute refresh and full source-to-destination lineage for all data streams. The built-in Dashboard focuses on marketing campaign metrics (opens, clicks, etc.) rather than end-to-end pipeline lineage. While Iterable's Smart Ingest feature tracks the status of data syncs and errors, it is not a visual lineage tool and runs at a minimum 15-minute interval (for scheduled syncs) rather than continuous sub-60s updates. In practice, ensuring data pipeline observability in Iterable would require external monitoring or custom integration, as Iterable doesn't natively visualize every first-party data flow in real time.</p>	https://support iterable.com/hc/en-us/articles/360015859471-Dashboard https://support iterable.com/hc/en-us/articles/24354008224916-Smart-Ingest-Overview
Yes	<p>2. Does your solution ship native inbound connectors or APIs for the top-10 commerce, marketing-automation, and service platforms without custom code or PS?</p> <p>Iterable provides multiple native integration options to ingest data from popular platforms. It has a robust REST API for custom integrations and ships pre-built connectors for many common tools. For example, Iterable's native integrations include connectors for e-commerce (Shopify via an app), CRM/service platforms (Zendesk bidirectional integration), and marketing data sources (Segment, Zapier, etc.). These allow teams to bring in data from top commerce, marketing, and support systems without heavy custom code. In addition, Iterable's Smart Ingest feature (co-developed with Hightouch) supports direct imports from data warehouses and storage (Redshift, BigQuery, S3, etc.). Overall, Iterable's ecosystem of inbound connectors and its APIs enable ingesting first-party customer data from a wide range of sources with minimal custom development.</p>	https://www.getcensus.com/research-blog-listing/Iterable-vs-emarsys-marketing-automation-powerhouses-compared https://support iterable.com/hc/en-us/articles/360021730551-Shopify-Iterable-Integration
Yes	<p>3. Does your solution support server-side and client-side collection for web, mobile-app, email, and in-store (POS/scanner) events in a single SDK?</p> <p>Iterable supports both client-side and server-side event collection across channels. It offers mobile SDKs for iOS, Android (and React Native) that automatically track in-app events, push opens, etc., and allow logging custom events from mobile apps. For web, Iterable can ingest events via its API or JavaScript (e.g., through integrations like the Shopify web pixel which captures site actions in real time). Email engagement events (sends, opens, clicks, unsubscribes) are tracked by Iterable's system automatically for campaigns. In-store/POS events can be sent server-to-server – for example, retailers can use Iterable's APIs to record purchase events or use Smart Ingest to sync offline transaction data into user profiles. All these events funnel into the unified profile, allowing a single customer view. In summary, Iterable's flexible SDKs and APIs enable data collection from web pages, mobile apps, emails, and offline sources in a unified manner (the “single SDK” concept is achieved via consistent APIs and data model across channels).</p>	https://support iterable.com/hc/en-us/articles/360028925511-Overview-of-Iterable-s-iOS-and-Android-SDKs https://support iterable.com/hc/en-us/articles/360021730551-Shopify-Iterable-Integration

Oct 2025 CDP Pulse Report - Iterable

Yes/No	Details	Source(s)
Yes	<p>4. Does your solution offer deterministic identity-resolution that stitches online cookies, MAIDs, loyalty IDs and POS receipts into one customer profile?</p> <p>Iterable uses deterministic keys to unify customer identities into one profile. Typically, an email address or a user-defined user ID serves as the unique identifier, and all data tied to that ID or email aggregates under one user profile. For example, Iterable recommends using email as the primary key to automatically merge activity across a user's devices (web, mobile app, etc.) into their profile. It also supports a custom userID for cases where email is not available; if a profile is initially identified by a userID (such as a loyalty or POS ID) in an email-based project, Iterable assigns a placeholder email to that profile to maintain uniqueness until a real email is provided. This deterministic approach means that as long as a common identifier (cookie ID associated with a userID/email, mobile advertising ID linked via profile data, loyalty number captured as a user field, etc.) is present, Iterable stitches the data into the same profile rather than creating duplicates. However, it does not perform fuzzy/probabilistic matching – identity resolution in Iterable relies on those explicit identifiers being passed or merged via API for a single customer view.</p>	https://support.iterable.com/hc/en-us/articles/360035402531-Identifying-the-User https://support.iterable.com/hc/en-us/articles/360035402531-Identifying-the-User
Yes	<p>5. Does your solution sync online and in-store inventory events in < 5 minutes (p95) to enable BOPIS and endless-aisle use cases?</p> <p>Iterable can handle near-real-time inventory updates via its API, though its scheduled batch syncs have a 15-minute floor. For example, if a product is bought in-store, the store system can call Iterable's API immediately to update that item's stock in the Iterable Catalog (product catalog) or trigger an event, achieving sub-5-minute propagation. Iterable's Smart Ingest scheduled syncs (often used for batch data from warehouses) run as frequently as every 15 minutes at best, but for truly immediate updates (like BOPIS inventory changes), developers are advised to use the real-time APIs or webhooks. In practice, this means you can stream inventory events (back-in-stock, low stock, etc.) to Iterable on an event-by-event basis. Iterable will process these incoming events virtually instantly – certainly fast enough for BOPIS/endless aisle scenarios where inventory visibility needs to update within a few minutes. While not a dedicated inventory module, Iterable's combination of API ingestion and profile/catalog updating supports <5min syncing when leveraged appropriately.</p>	https://support.ometria.com/hc/en-gb/articles/57324428317-Setting-up-an-API-based-Magento-2-integration
No	<p>6. Does your solution expose decision or personalization APIs that return in <1 second (p95) so actions fire while the session is active?</p> <p>Iterable does not provide a standalone real-time decision engine API (for example, a next-best-action or personalization lookup API) that external applications can call for sub-second responses. Its platform personalizes content within campaigns and journeys – e.g., using stored profile data or precomputed recommendations – but there's no documented public API that performs on-demand decisioning with <1s latency. Iterable's API endpoints are primarily for data CRUD operations (managing users, triggering sends, querying metrics), rather than dynamic in-session personalization. Marketers can achieve real-time triggers within Iterable (for instance, an event can immediately trigger a message or workflow while a user is on site), but if a web app wanted to query Iterable for a personalization decision during a live session, there isn't an out-of-the-box inference service. In summary, Iterable's focus is on real-time messaging execution (e.g., sending an email or push instantly in response to behavior) rather than exposing a millisecond-level personalization API for arbitrary downstream use.</p>	https://ometria.com/use-cases/experience-orchestration https://docs.ometria.com/docs/getting-started-with-website-personalisation

Oct 2025 CDP Pulse Report - Iterable

Yes/No	Details	Source(s)	
7. Does your solution include embedded predictive-merchandising models (e.g., next-best product) deployable by business users without data-science help?	No	<p>Iterable provides some AI-driven features (such as send-time optimization and Predictive Goals for propensity scoring), but it does not offer built-in product recommendation or “next-best-product” models that a marketer can deploy out-of-the-box. Unlike some CDPs or marketing clouds with one-click predictive merchandising, Iterable’s AI Suite is oriented toward optimizing messaging (when to send, subject line experimentation, etc.) rather than automatic product recommendations. For merchandising use cases, teams would typically have to build their own recommendation logic externally (or use a partner integration) and then import those results into Iterable’s Catalog or user profiles for personalization. The platform’s native capabilities around predictive analytics are more limited in depth compared to specialized retail marketing. Therefore, business users cannot simply turn on a built-in “next best product” model in Iterable – any such model would require custom data science work outside the platform, with the outputs fed into Iterable for campaign use.</p>	https://www.getcensus.com/research-blog-listing/iterable-vs-emarsys-marketing-automation-powerhouses-compared https://www.getcensus.com/research-blog-listing/iterable-vs-emarsys-marketing-automation-powerhouses-compared
8. Does your solution allow teams to upload and deploy their own ML/AI models (PMML/ONNX/Python) to score profiles or events in real time?	No	<p>Iterable does not support uploading custom ML models for execution within the platform. It has no feature to import PMML/ONNX models or custom Python code and then run those models on profile/event data in real time. Iterable’s AI capabilities are limited to the built-in features it provides (e.g., predictive goal scoring done by Iterable’s internal algorithms, and some content-generation AI). There is no self-serve “bring your own model” functionality. If a team wants to apply a proprietary machine learning model (for example, a custom churn score or product affinity model), they would need to score users externally (in their own environment or a data science platform) and then import the results (e.g., as user profile fields or segments) via Iterable’s APIs or batch uploads. In short, Iterable cannot host or execute arbitrary user-provided ML models in real-time – any custom model scoring must be done outside and then integrated into Iterable’s data.</p>	https://www.getcensus.com/research-blog-listing/iterable-vs-emarsys-marketing-automation-powerhouses-compared https://www.getcensus.com/research-blog-listing/iterable-vs-emarsys-marketing-automation-powerhouses-compared
9. Does your solution provide an SQL workspace or direct lake-house access for segmentation, attribution, and ad-hoc analytics?	No	<p>Iterable does not offer a built-in SQL query editor or direct access to a cloud data lake/warehouse for customer data. Segmentation and analytics in Iterable are handled through its UI (segmentation builder, pre-built reports) rather than raw SQL queries. There isn’t a feature where analysts can write arbitrary SQL against Iterable’s data store. Instead, Iterable encourages exporting data to external systems for advanced analysis – for example, some users leverage third-party tools like Census to sync Iterable data with data warehouses for complex queries. Iterable’s native data integration rating is modest (it’s noted that data warehouse support is better when using an add-on tool) getcensus.com. While you can retrieve data via APIs or set up system webhooks to pipe events outward, the platform itself doesn’t provide a direct Snowflake/BigQuery access or SQL workspace. Thus, ad-hoc analytics or multi-touch attribution often require moving Iterable data into an external BI environment or using Iterable’s preset reports rather than writing your own SQL within Iterable.</p>	https://www.getcensus.com/research-blog-listing/iterable-vs-emarsys-marketing-automation-powerhouses-compared https://www.getcensus.com/research-blog-listing/iterable-vs-emarsys-marketing-automation-powerhouses-compared

Oct 2025 CDP Pulse Report - Iterable

Yes/No	Details	Source(s)
No	<p>Iterable does not have an automated data observability module that performs anomaly detection on incoming data or quarantines “bad” data. The platform will validate and error-out individual records that don’t meet schema requirements (for example, if a field is badly formatted or an event payload is invalid, Iterable’s import process will mark it as an error or warning). However, it doesn’t proactively learn normal data patterns or alert on anomalies in metrics like event volume spikes/dips. Data quality controls in Iterable are mostly limited to field mappings and basic validations (and its system webhooks or alerts for failures in integrations). There is no built-in ML-driven monitoring that quarantines data automatically. Ensuring upstream data quality and building anomaly detection would largely be up to the user’s engineering/BI practices outside of Iterable. In summary, while Iterable will log errors for bad data (and allow you to see failures in Smart Ingest or API calls), it doesn’t offer a full-fledged continuous data-quality governance or anomaly detection feature out-of-the-box.</p>	https://support.iterable.com/hc/en-us/articles/25908941582996-Monitoring-Smart-Ingest-Syncs https://support.iterable.com/hc/en-us/articles/25908941582996-Monitoring-Smart-Ingest-Syncs
Yes	<p>Iterable includes enterprise-grade security features such as role-based access control (RBAC) and auditing, although field-level granularity is limited. RBAC: Access to data and actions in Iterable is managed by user roles/permissions – data is only accessible after proper authentication, and permissions can be configured by role (e.g., who can download data, who can manage integrations, etc.). Audit logs: Iterable maintains logs of key account activities (for compliance, Iterable is SOC 2 certified), though these logs are not exposed for every single field change to customers in the UI. (They provide admin-facing audit trails for changes like user access, and one can request logs for compliance needs.) Data residency: All these controls apply in each region – Iterable offers separate regional deployments (US and Europe) to keep data in-region as required, with the same security controls (encryption, RBAC) in both. In practice, Iterable allows fine-grained project-level or feature-level permissions, but not an out-of-the-box “immutable field-level change log” visible to end users. Nonetheless, its security model ensures that across regions, access is restricted by role and any configuration changes can be audited internally.</p>	https:// iterable.com/features/centralized-data/ https://support.iterable.com/hc/en-us/articles/17527250887444-Iterable-s-European-Data-Center-EDC
Yes	<p>Iterable has built-in support for consent and preference management. It uses message channels and message types to categorize communications, allowing users to opt out at a granular level (for example, unsubscribe from SMS but not email, or from a specific list/newsletter). The platform provides hosted default unsubscribe pages out-of-the-box, and also APIs to build custom preference centers. When a user updates their consent/preferences, Iterable immediately reflects that – e.g., if a user clicks an unsubscribe link, that change is logged and the user will be suppressed from further sends on that channel/type essentially instantly (no additional manual steps). Iterable’s documentation emphasizes timely processing of unsubscribe requests to remain compliant. Additionally, Iterable supports GDPR compliance by providing tools to handle data subject requests (like right to be forgotten) and by segregating EU user data in the European data center when needed. In summary, Iterable is designed to honor consent changes in real time, preventing any non-compliant activations, and offers the infrastructure for preference management required for regulations like GDPR/CCPA.</p>	https://www.getcensus.com/research-blog-listing/Iterable-vs-emarsys-marketing-automation-powerhouses-compared https://www.getcensus.com/research-blog-listing/Iterable-vs-emarsys-marketing-automation-powerhouses-compared

Oct 2025 CDP Pulse Report - Iterable

Yes/No	Details	Source(s)	
13. Does your solution support regional data residency options (e.g., EU-only, US-only) with encryption at rest (AES-256) and in transit (TLS 1.2+)?	Yes	<p>Iterable supports regional data residency and strong encryption standards. It offers a European Data Center (EDC) option for customers who need EU-only data storage – projects can be hosted on the EU (Ireland) data center or the US data center as contracted, ensuring that personal data stays within the chosen region. Iterable confirms that the EU environment is GDPR-compliant and does not send data to the US, with feature parity to the US environment. In terms of security, all data in Iterable is encrypted in transit and at rest. Iterable's security page notes that data is encrypted when moving through the system and when stored on disk. (While specific algorithms aren't quoted on the public docs, the industry-standard AES-256 at rest and TLS 1.2+ for transit can be inferred from compliance certifications.) Additionally, access to data is controlled by authentication and roles, as part of its enterprise security posture. So yes – Iterable meets regional residency needs and employs modern encryption to protect data both at rest and in motion.</p>	https://www.getcensus.com/research-blog-listing/Iterable-vs-emarsys-marketing-automation-powerhouses-compared https://www.getcensus.com/research-blog-listing/Iterable-vs-emarsys-marketing-automation-powerhouses-compared
14. Does your solution provide reverse-ETL / audience-syndication connectors that push profiles and events to major ad, analytics, and messaging destinations natively?	Yes	<p>Iterable can act as a source of data to other systems through native integrations and webhooks. For example, Iterable has a built-in Facebook Custom Audiences integration that lets you export user lists directly from Iterable to Facebook Ads as custom audience segments. Marketers can automate audience sync to Facebook, including using a Journey step to add users to or remove from a Facebook audience in real-time. Similarly, Iterable data can be pushed to other ad or analytics platforms using system webhooks or partner connectors – e.g., sending conversion events to Google Analytics or forwarding audiences to Google Ads can be achieved via Iterable's integrations (Zapier or API).</p>	https://www.getcensus.com/research-blog-listing/Iterable-vs-emarsys-marketing-automation-powerhouses-compared https://www.getcensus.com/research-blog-listing/Iterable-vs-emarsys-marketing-automation-powerhouses-compared
15. Does your solution expose a unified customer-profile API/stream that any downstream system can query or subscribe to in real time?	Yes	<p>Iterable offers a unified profile API and event streaming capabilities for real-time access. Every user in Iterable has a single consolidated profile that can be retrieved on-demand via the API – for example, developers can use the GET /api/users/{userId}/email endpoint to fetch the latest profile attributes for a given customer. This allows any downstream system to query a customer profile in real time (for instance, to see attributes or segment membership before taking an action). Additionally, Iterable's system webhooks enable a push-based stream of events: you can subscribe an external endpoint to receive real-time data whenever certain events occur (such as profile updates like subscribes/unsubscribes, message sends, opens, clicks, etc.). While Iterable's webhooks focus on event streams rather than a change-data-capture of every profile attribute, they effectively let downstream systems know what's happening with users in real time. In summary, Iterable provides both a pull API for unified profiles and a push mechanism (webhooks) for live event streaming, so external systems can query or listen and react to customer data in near real time.</p>	https://support.iterable.com/hc/en-us/articles/208463956-Creating-a-Subscription-Preference-Center https://support.iterable.com/hc/en-us/articles/208013936-System-Webhooks
16. Does your solution publish transparent, usage-based pricing (events, MTU, storage) in writing, independent of professional-services spend?	No	<p>Iterable's pricing is not published transparently on a usage-based rate card. The company uses a custom pricing model based primarily on the number of contacts (users) and the feature package, typically structured in tiers (e.g., Growth, Scale, Enterprise). They do not publicly list prices per event, MAU/MTU, or GB of storage on their website. Instead, pricing is provided via quotes and can vary by customer needs – for instance, reports indicate a starting price around \$3,000 per month for the base tier, but exact costs scale with contact counts and messaging volume. There is no self-serve pricing calculator detailing usage-based costs independent of services. Professional services (like onboarding or deliverability consulting) are separate, but the core platform pricing itself isn't transparently usage-metered in marketing materials. In summary, Iterable does not offer openly published usage-based pricing documentation; prospective customers must engage with sales for a tailored quote rather than refer to a public price list.</p>	https://www.getcensus.com/research-blog-listing/Iterable-vs-emarsys-marketing-automation-powerhouses-compared https://www.getcensus.com/research-blog-listing/Iterable-vs-emarsys-marketing-automation-powerhouses-compared

Vendor Deep Dive



SAP Emarsys

Primary Archetype: Commerce & Personalization Engine

Strengths: Retail Specialist; Real-time/Edge

Capability Lens: In-Session Activation

VOX Score: 81.06 / 100

Preference Coverage: 13 / 16 questions “Yes”

📞 Talk to SAP Emarsys: <https://emarsys.com/customer-engagement-platform/>

📞 [Ask VOX for Introduction](#)

Oct 2025 CDP Pulse Report - SAP Emarsys

Yes/No	Details	Source(s)	
1. Does your solution provide data-observability dashboards that refresh in ≤ 60 seconds (p95) and show source-to-destination lineage for every first-party data stream?	Yes	Emarsys offers data integration monitoring dashboards. For example, the "Predict Data Sources" page provides a status overview of data feeds (web behavior, product catalog, sales data) and flags any integration issues in near real time, helping ensure first-party data flows are healthy. These tools let users validate that data from sources (e.g. website, e-commerce) is being received and processed almost instantly (within seconds) by Emarsys. While it may not render a full graphical lineage, it gives marketers rapid visibility into each data stream's status and continuity.	https://training.emarsys.com/hc/en-us/articles/360014444658-HowTo-Video-Web-Extend-Predict-Data-Sources
2. Does your solution ship native inbound connectors or APIs for the top-10 commerce, marketing-automation, and service platforms without custom code or PS?	Yes	Emarsys provides many out-of-the-box connectors for popular commerce and customer systems. It has plug-and-play integrations for leading e-commerce platforms (e.g. SAP Commerce Cloud, Shopify Plus, Adobe Commerce/Magento, Salesforce Data 360, BigCommerce), enabling automatic ingestion of customers, products, orders and events without custom development. Emarsys also natively integrates with SAP's CX suite (Sales Cloud, Customer Data Cloud, etc.) and offers open APIs, so brands can connect marketing data from top CRM, service, and loyalty systems without heavy coding or professional services.	https://emarsys.com/emarsys-sap/ https://emarsys.com/partner-ecosystem/directory/technology-partners/big-commerce-sap-emarsys-plugin/
3. Does your solution support server-side and client-side collection for web, mobile-app, email, and in-store (POS/scanner) events in a single SDK?	Yes	Emarsys supports multi-channel data collection across client and server. For web, it uses a client-side JavaScript SDK (Web Extend) to track page views and behavior in browsers. For mobile apps, Emarsys provides iOS/Android SDKs to collect in-app events and enable push/in-app messaging. These SDKs feed data into the same platform. Additionally, Emarsys's server-side APIs capture events like purchases or POS transactions from back-end systems. For example, its BigCommerce plugin streams web interactions, product catalog updates, inventory changes, carts, and orders into Emarsys. All these first-party events (online and offline) converge into a unified customer data layer for real-time use.	https://help.emarsys.com/hc/en-us/articles/115002683889-Overview-Mobile-Engage-Overview https://emarsys.com/partner-ecosystem/directory/technology-partners/big-commerce-sap-emarsys-plugin/
4. Does your solution offer deterministic identity-resolution that stitches online cookies, MAIDs, loyalty IDs and POS receipts into one customer profile?	Yes	Emarsys creates a single unified customer profile by deterministically matching identifiers. Anonymous web visitors tracked by cookies or mobile device IDs are linked to known contacts once they log in or provide an email, stitching their prior behavior to their profile. Emarsys's data model merges data from loyalty programs, in-store POS transactions (via email/ID capture at checkout), and online interactions into one record. This "single customer view" ensures that a customer's online and offline identities (e.g. loyalty ID, email, POS receipt info) are tied together for personalization.	https://emarsys.com/learn/blog/email-digital-id/ https://emarsys.com/partner-ecosystem/directory/technology-partners/salesforce-commerce/
5. Does your solution sync online and in-store inventory events in < 5 minutes (p95) to enable BOPIS and endless-aisle use cases?	Yes	Emarsys can reflect inventory changes from stores and online in near real-time via its commerce integrations. It ingests product and stock updates through native connectors (e.g. the BigCommerce-Emarsys plugin) syncs product catalog and inventory events continuously. This enables scenarios like Buy-Online-Pickup-In-Store (BOPIS) and endless aisle: for example, when a customer places an order for in-store pickup, Emarsys can immediately trigger communications and ensure recommendations exclude out-of-stock items. Emarsys even supports use cases to upsell BOPIS customers with items currently in-stock at that pickup store, indicating that inventory data is up-to-date within minutes.	https://emarsys.com/partner-ecosystem/directory/technology-partners/big-commerce-sap-emarsys-plugin/ https://emarsys.com/cross-channel-marketing-automation/instore/
6. Does your solution expose decision or personalization APIs that return in < 1 second (p95) so actions fire while the session is active?	Yes	Emarsys provides real-time personalization capabilities via APIs/SDKs to drive in-session actions. For instance, its product recommendation engine (part of Emarsys Predict) delivers content instantly on websites or apps – when a user is browsing, Emarsys can return personalized product suggestions or web content within sub-second latency. The platform's Open-time personalization for emails and the Web Channel personalization ensure responses fast enough for live interactions. Emarsys also supports real-time decisioning in customer journeys: if a user triggers an event (like form fill or add-to-cart), the Emarsys Interactions service will react immediately (typically well under 1s) to fire off the next best action (such as an on-site message or email) with no delay. These swift response times allow personalization to occur while the customer is still engaged in the session.	https://emarsys.com/learn/blog/emarsys-spring-release-2023/ https://emarsys.com/partner-ecosystem/directory/technology-partners/big-commerce-sap-emarsys-plugin/
7. Does your solution include embedded predictive-merchandising models (e.g., next-best product) deployable by business users without data-science help?	Yes	Emarsys comes with built-in AI predictive models for merchandising that marketers can use out-of-the-box. The platform's AI scores each customer for things like "next-best purchase" product, likelihood to engage, or churn risk. These predictive recommendations (e.g. suggesting the next product a customer is likely to buy) automatically feed into personalization features – populating emails, web recommendations, or segments – without requiring any data science work. Emarsys's predictive analytics (e.g. product affinity and "also bought" algorithms) are embedded and tuned within the solution, so business users can deploy next-best product recommendations and similar models via simple UI settings, rather than coding ML models.	https://emarsys.com/learn/blog/predictive-customer-analytics/ https://emarsys.com/learn/blog/predictive-customer-analytics/

Oct 2025 CDP Pulse Report - SAP Emarsys

	Yes/No	Details	Source(s)
8. Does your solution allow teams to upload and deploy their own ML/AI models (PMML/ONNX/Python) to score profiles or events in real time?	No	Emarsys does not support uploading custom external machine-learning models for deployment in the platform. Instead, Emarsys focuses on its native AI capabilities (churn models, product recommendations, etc.) which are pre-integrated. Users cannot import a PMML/ONNX model or custom Python ML code to run scoring within Emarsys. The platform's AI is leveraged through built-in features – marketers use the provided models and predictions (which Emarsys continuously maintains) rather than deploying their own algorithms in the system. Any advanced custom modeling would need to be done outside of Emarsys and then fed in via data import or API, as Emarsys's environment doesn't host user-supplied ML models.	https://emarsys.com/learn/blog/predictive-customer-analytics/
9. Does your solution provide an SQL workspace or direct lake-house access for segmentation, attribution, and ad-hoc analytics?	Yes	Emarsys offers an Open Data solution that gives customers direct access to their raw data for analysis. This feature (built on Google BigQuery) lets users query email, web, mobile, and campaign data using SQL and even connect BI tools without manual exports. With Open Data, marketers or analysts can perform ad-hoc segmentation, attribution modeling, and custom analytics by running SQL queries on engagement data in a cloud data warehouse environment. In addition to this, Emarsys provides an API and UI for standard segmentation and reporting, but the Open Data workspace is available for deeper analytical needs, effectively functioning as a direct lake-house access to the customer and event data.	https://www.aplytosupply.digitalmarketplace.service.gov.uk/g-cloud/services/835807560526235
10. Does your solution run continuous data-quality and anomaly detection with automatic quarantine or alerting before bad data feeds downstream systems?	No	Emarsys performs basic data validation on import and provides health dashboards, but it does not have autonomous anomaly detection or quarantining of data anomalies before downstream use. The platform's data load monitors (e.g. for contact, product, or sales data feeds) will flag errors like format mismatches or missing fields and can send error alerts to users, but there isn't a full AI-driven anomaly detection system evaluating data quality trends. Users must review the status dashboards (such as the Smart Insight Data Quality screens) for issues. If bad data is detected (e.g. via validation rules), Emarsys will skip or error those records and log the issue, but it won't automatically quarantine data streams – instead it relies on admins to correct issues once alerted.	https://training.emarsys.com/hc/en-us/articles/360014444658-HowTo-Video-Web-Extend-Predict-Data-Sources
11. Does your solution enforce role-based access controls and immutable audit logs at the field level across all regions?	Yes	Emarsys has robust role-based access control (RBAC) and auditing features. Administrators can define user roles and granular permissions – down to which features or data areas each user can access – ensuring, for example, that a user can be restricted from viewing or editing specific sensitive contact fields or pages. All user actions are tracked for compliance: Emarsys maintains audit logs of changes (e.g. data imports, campaign launches, permission changes). These logs are immutable and audit-ready, helping meet GDPR/CCPA requirements. In practice, this means every modification to customer data or settings is recorded and can be reviewed. Emarsys's commitment to security and compliance is reflected in their approach to field-level encryption and audit trails which are consistent across regions (EU, US, etc.), so data governance standards remain uniformemarsys.com.	https://www.aplytosupply.digitalmarketplace.service.gov.uk/g-cloud/services/835807560526235 https://emarsys.com/learn/blog/how-ryderwear-uses-predictive-analytics-marketing-to-unlock-success/
12. Does your solution offer built-in consent & preference management that suppresses activation instantly when permissions change (GDPR/CCPA ready)?	Yes	Emarsys includes native consent and preference management. It tracks consent status for each contact (such as email opt-in, SMS permission, privacy consent flags) and honors any change immediately across all channels. If a customer withdraws consent or opts out (for example, unsubscribing from emails or changing cookie preferences), Emarsys instantly reflects that preference and will suppress any further messaging to that individual until they opt back in. This instantaneous enforcement ensures compliance with GDPR/CCPA – no communications go out if permission is revoked. Emarsys also provides tools like subscription centers and APIs to capture consent updates, and it logs those changes so that consent is centralized and audit-ready for regulators.	https://emarsys.com/learn/blog/how-ryderwear-uses-predictive-analytics-marketing-to-unlock-success/
13. Does your solution support regional data residency options (e.g., EU-only, US-only) with encryption at rest (AES-256) and in transit (TLS 1.2+)?	Yes	Emarsys allows clients to keep data in specific regions to meet data residency requirements. For example, EU customer data can be hosted on EU-only servers (within the EEA) to comply with regulations, and similarly Emarsys offers US or other regional hosting options. Security-wise, Emarsys enforces strong encryption: all customer data at rest is encrypted (using industry-standard AES-256 encryption on storage media) and data in transit is protected via TLS 1.2+ secure protocols. These measures ensure that whether data is stored or being transferred, it remains encrypted and secure. Emarsys's cloud infrastructure meets enterprise security certifications, and clients have transparency and control regarding where their data is stored geographically and how it's protected.	https://www.aplytosupply.digitalmarketplace.service.gov.uk/g-cloud/services/835807560526235 https://www.aplytosupply.digitalmarketplace.service.gov.uk/g-cloud/services/835807560526235

Oct 2025 CDP Pulse Report - SAP Emarsys

	Yes/No	Details	Source(s)
14. Does your solution provide reverse-ETL / audience-syndication connectors that push profiles and events to major ad, analytics, and messaging destinations natively?	Yes	<p>Emarsys can natively syndicate audiences and event data to a variety of external marketing and analytics destinations. It has built-in connectors to push customer segments (profiles) to major ad platforms like Google Ads, Facebook/Meta (Instagram), LinkedIn, TikTok, and Criteo, enabling marketers to create matched audiences or retargeting campaigns directly from Emarsys data. Emarsys also integrates with external messaging services – for instance, an SMS Integration API allows using third-party SMS providers instantly with Emarsys campaigns. For analytics, Emarsys can feed data into tools like Google Analytics or BigQuery (via its Open Data export). These “reverse ETL” style features mean Emarsys doesn’t just ingest data – it actively distributes customer data and campaign events out to other marketing ecosystems in a turnkey way, without requiring custom coding.</p>	https://emarsys.com/omnichannel-integrations/advertising-integrations/ https://emarsys.com/omnichannel-integrations/advertising-integrations/
15. Does your solution expose a unified customer-profile API/stream that any downstream system can query or subscribe to in real time?	Yes	<p>Emarsys exposes its unified customer profiles through APIs, allowing external systems to query profile data in real time. The Emarsys Core API lets you retrieve and update contact records (including all profile fields, segmentation data, opt-ins, etc.) on demand. This means any authorized downstream application (such as a web portal, CRM, or custom app) can call Emarsys to get the latest customer profile details (for example, to personalize content or check preference flags). While Emarsys’s API is request/response (pull) based, it effectively serves as a unified customer profile service. Additionally, for event streaming, Emarsys offers webhook-like capabilities (e.g., triggering external calls via External Events), but it primarily relies on the real-time API access to empower other systems with up-to-date profile information.</p>	https://www.aplytosupply.digitalmarketplace.service.gov.uk/g-cloud/services/835807560526235
16. Does your solution publish transparent, usage-based pricing (events, MTU, storage) in writing, independent of professional-services spend?	No	<p>Emarsys does not publicly disclose a transparent usage-based pricing model on its website. Its pricing is typically obtained via sales consultation and is generally a tiered or custom enterprise model. In practice, Emarsys pricing is based on factors like number of contacts (or monthly active contacts) and desired product modules, rather than openly posted metered rates for events or storage. Independent of services, the costs are bundled and quoted per contract. Industry comparisons note that Emarsys is a premium solution and that it does not provide public, usage-based pricing information, with enterprise licenses often negotiated on a case-by-case basis. Prospective customers must contact Emarsys for a detailed quote; the pricing is not transparently published in a self-serve manner.</p>	https://www.getcensus.com/research-blog-listing/emarsys-vs-braze-comparing-top-marketing-automation-platforms

Vendor Deep Dive



Bloomreach

Primary Archetype: Commerce & Personalization Engine

Strengths: Engagement Hub; Retail Specialist;
Real-time/Edge

Capability Lens: In-Session Activation

VOX Score: 81.12 / 100

Preference Coverage: 13 / 16 questions “Yes”

📞 Talk to Bloomreach: <https://www.bloomreach.com/>

📞 [Ask VOX for Introduction](#)

Oct 2025 CDP Pulse Report - Bloomreach

Yes/No	Details	Source(s)
No	Bloomreach Engagement offers real-time analytics dashboards for monitoring key metrics, but it does not explicitly mention 60-second refresh guarantees or a dedicated source-to-destination lineage visualization for all data streams in its documentation. Dashboards update with near real-time insights (e.g. conversion funnels use real-time data) but we found no evidence of a full data lineage UI within the platform.	https://documentation.bloomreach.com/engagement/docs/dashboards-1 https://www.bloomreach.com/en/blog/bloomreach-engagement-custom-notification
Yes	Bloomreach provides out-of-the-box connectors for many popular commerce and marketing platforms. For example, it has native integrations for Shopify, Magento, BigCommerce, etc., and built-in connectors for marketing and advertising channels (Facebook/Instagram Ads, Google Ads, email/SMS providers, etc.) that can be enabled without custom coding. The documentation shows a variety of third-party integrations available, and if a system isn't pre-listed, Bloomreach offers open APIs to connect virtually any platform.	https://documentation.bloomreach.com/engagement/docs/shopify-and-shopify-plus-overview https://documentation.bloomreach.com/engagement/docs/integrations
Yes	Bloomreach Engagement can ingest data from client-side sources (web JavaScript SDK, mobile SDKs for iOS/Android) as well as server-side sources. It tracks website behavior, mobile app events, email interactions (opens/clicks), and even offline or in-store events (via its APIs or batch imports). All these events funnel into the same unified customer profile. In practice, Bloomreach's Customer Data Engine unified purchase data (including in-store transactions), web browsing, and mobile app engagement data in one place. This means both client and server events are supported and merged into a single view of the customer.	https://www.imdigital.com/blog/flybox-cdp-marketing-automation https://documentation.bloomreach.com/engagement/docs/customer-identification
Yes	Bloomreach Engagement uses deterministic identity resolution via configurable hard IDs and soft IDs. For example, an anonymous user tracked by a browser cookie (soft ID) can be linked to an identified profile once they log in or make a purchase (hard ID like email or loyalty ID). The platform allows multiple identifiers (e.g. email, loyalty card number, mobile device ID) on one profile and automatically merges events accordingly. Bloomreach's docs show that a new visitor starts with a cookie ID and then, upon registration or purchase, that activity is tied to their email/CRM ID, unifying past anonymous behavior with the known profile. It also supports additional IDs like offline loyalty IDs or POS transaction IDs as hard identifiers to stitch data into the single customer view.	https://documentation.bloomreach.com/engagement/docs/customer-identification https://documentation.bloomreach.com/engagement/docs/customer-identification
Yes	Bloomreach can handle near real-time synchronization of product and inventory data between online and offline systems. For example, the Shopify integration continuously imports product catalog updates "in real time" and listens for updates on orders/customers. This implies inventory changes on the e-commerce side are pulled in promptly. While an exact "<5 minute" SLA isn't explicitly stated, the system is designed for fast data updates. This real-time product data capability enables use cases like Buy Online, Pickup In-Store and endless aisle by ensuring that inventory information from stores and online channels stays up-to-date within minutes.	https://documentation.bloomreach.com/engagement/docs/shopify-and-shopify-plus-overview
Yes	Bloomreach offers high-performance personalization and decisioning via its APIs. In practice, the platform's AI-driven recommendation engine can serve personalized product recommendations "instantaneously (within 0.1 seconds of a user's action)". This sub-second response capability means that when a user is active on a site or app, Bloomreach's decisioning (e.g. next-best product recommendations, content personalization API calls) can deliver results fast enough to be inserted into the current page or session. The system is built on an in-memory framework for split-second data access, enabling API responses well under 1 second for real-time personalization.	https://www.bloomreach.com/en/blog/what-is-real-time-personalization https://www.bloomreach.com/en/blog/what-is-real-time-personalization
Yes	Bloomreach Engagement comes with built-in predictive analytics and merchandising models powered by its Loomi AI. Non-technical users (marketers/merchandisers) can deploy these via a library of pre-built templates and a point-and-click interface. For example, the platform provides prediction templates (for churn risk, purchase likelihood, optimal send time, etc.) and product recommendation models out of the box. Users can configure and launch these models (such as next-best product recommendations or propensity scores) without writing code or needing a data scientist – the system auto-trains and applies them. The documentation highlights an "easy-to-use predictions builder" and ready-to-use AI templates for common use cases, confirming that business users can leverage predictive merchandising models on their own.	https://documentation.bloomreach.com/engagement/docs/predictions-generally https://documentation.bloomreach.com/engagement/docs/product-recommendations

Oct 2025 CDP Pulse Report - Bloomreach

Yes/No	Details	Source(s)
8. Does your solution allow teams to upload and deploy their own ML/AI models (PMML/ONNX/Python) to score profiles or events in real time?	No Bloomreach's approach to AI is to provide its own built-in machine learning (Loomi AI) capabilities, rather than letting clients import custom model files. We did not find any feature that allows uploading a PMML or ONNX model for deployment within Bloomreach. The platform's Predictions engine is powered by Bloomreach's native algorithms (Loomi) and offers "custom predictions" in-app (which are configured through the UI, not by importing external code). There is no indication that a team can bring their own arbitrary Python/ONNX model to run inside Bloomreach in real-time. Any custom modeling would likely have to be done outside the platform and integrated via API, as Bloomreach does not natively host user-provided ML models.	https://documentation.bloomreach.com/engagement/docs/predictions-manual https://documentation.bloomreach.com/engagement/docs/predictions-generally
9. Does your solution provide an SQL workspace or direct lake-house access for segmentation, attribution, and ad-hoc analytics?	Yes Bloomreach Engagement offers robust options for advanced analytical access. Notably, it provides Engagement BigQuery (EBQ) – a fully managed BigQuery data warehouse instance that is populated with the customer's Bloomreach. This gives analysts direct SQL access to all their customer and event data for ad-hoc queries, building segments, custom attribution models, etc. Users can connect BI tools or use a SQL workspace on this data. Additionally, within the platform, Bloomreach has a "SQL Reports" feature for running custom queries on the data. Together, these capabilities mean teams can either write SQL in-platform or query the exported data in a lakehouse environment, enabling deep analytical flexibility.	https://cloud.google.com/blog/products/data-analytics/bloomreach-engagement-integrates-with-bigquery/ https://cloud.google.com/blog/products/data-analytics/bloomreach-engagement-integrates-with-bigquery/
10. Does your solution run continuous data-quality and anomaly detection with automatic quarantine or alerting before bad data feeds downstream systems?	Yes Bloomreach is adding AI-powered anomaly detection to proactively flag data issues. According to Bloomreach, the platform will use statistical techniques to monitor metrics and will "send notifications when an anomaly is detected," allowing users to address data quality problems early. This feature is designed to catch unusual patterns (e.g. sudden drops or spikes in events) in real time. While the documentation (as of 2023) describes this as an emerging feature, it indicates a commitment to continuous data-quality monitoring. The system does not explicitly mention quarantining bad data automatically, but it does support alerting so that suspect data can be investigated or stopped before it impacts downstream processes.	https://www.bloomreach.com/en/blog/bloomreach-engagement-custom-notification https://www.bloomreach.com/en/blog/bloomreach-engagement-custom-notification
11. Does your solution enforce role-based access controls and immutable audit logs at the field level across all regions?	Yes Bloomreach Engagement has robust security controls including fine-grained RBAC and audit logging. Administrators can define roles that restrict access to specific data – for example, there is a special "Personal Data Viewer" role required to view sensitive personal fields (PII). This ensures field-level protection of data like customer identifiers across all users/regions. Additionally, Bloomreach maintains an audit log of user activities in the platform. The audit logs are stored in a tamper-proof manner (exported to a secure Google Cloud Storage bucket with read-only access) so they cannot be altered. This immutable audit trail records who did what and when, which is crucial for compliance and security audits.	https://documentation.bloomreach.com/engagement/docs/access-management https://documentation.bloomreach.com/engagement/docs/audit-log
12. Does your solution offer built-in consent & preference management that suppresses activation instantly when permissions change (GDPR/CCPA ready)?	Yes Bloomreach Engagement includes integrated consent and preference management. Users can define multiple consent categories (e.g. Email Marketing, SMS, Tracking) and manage customer opt-ins/opt-outs per category. The platform uses these settings in real time: when you execute campaigns or journeys, you must select the applicable consent category, and Bloomreach will only send to customers who have given that consent. If a customer withdraws consent, the change takes effect immediately – for example, they will be automatically excluded from any future sends in that category. This built-in mechanism supports GDPR/CCPA compliance by instantly honoring permission changes and not requiring any manual intervention to suppress communications. Tracking consent for website data collection is also supported, with options to require explicit opt-in before tracking, ensuring privacy regulations are met.	https://documentation.bloomreach.com/engagement/docs/consent-categories https://documentation.bloomreach.com/engagement/docs/consent-categories

Oct 2025 CDP Pulse Report - Bloomreach

	Yes/No	Details	Source(s)
13. Does your solution support regional data residency options (e.g., EU-only, US-only) with encryption at rest (AES-256) and in transit (TLS 1.2+)?	Yes	Bloomreach allows customers to host data in regional environments to meet data residency requirements. (For instance, Bloomreach's cloud infrastructure can be set up in EU or US data centers based on client preference.) All data is secured via strong encryption: encryption at rest is enabled by default using AES-256 (managed through Google Cloud's encryption services), and all data in transit is protected with TLS 1.2+ secure protocols. The platform emphasizes security-by-design – it even notes that by default "data is encrypted both at rest and in transit". Combined with certifications like ISO 27001 and SOC2, Bloomreach meets enterprise security and privacy standards in every region it operates.	https://documentation.bloomreach.com/engagement/docs/security-architecture https://www.bloomreach.com/en/security-at-bloomreach
14. Does your solution provide reverse-ETL / audience-syndication connectors that push profiles and events to major ad, analytics, and messaging destinations natively?	Yes	Bloomreach Engagement has native audience syndication capabilities through its integration library. It can directly push customer profiles, segments, or events to a variety of downstream systems. For example, it includes built-in connectors for Facebook/Instagram Ads, Google Ads (Customer Match), and even a Facebook Conversions API – allowing marketers to automatically sync audiences for ad targeting. Similarly, it can send data to analytics platforms or data warehouses (BigQuery, Google Analytics) and to messaging services (like via webhooks or built-in email/SMS integrations). The documentation explicitly mentions that integrations are available for sending data to Facebook, Google, AdForm and more without additional custom development. These out-of-the-box reverse-ETL connectors enable Bloomreach to act as a central hub, syndicating audience data to all major marketing and advertising endpoints.	https://documentation.bloomreach.com/engagement/docs/integrations https://documentation.bloomreach.com/engagement/docs/integrations
15. Does your solution expose a unified customer-profile API/stream that any downstream system can query or subscribe to in real time?	Yes	Bloomreach provides APIs to access the unified customer profile data on demand. It offers a Customer API that lets you retrieve a full profile (all attributes and even event history) for a given customer ID. In fact, there is an "Export a customer" endpoint which "retrieves all customer attributes and events...for a single customer." This means external systems can query Bloomreach in real time to get the latest state of a customer profile (for example, to personalize an external application or service). While Bloomreach doesn't stream profiles via a push subscription out-of-the-box, it supports real-time pulls and can also send event updates via webhooks if configured. Essentially, any downstream system can call the Bloomreach API to get a Single Customer View at any time, making the rich profile data accessible outside the CDP.	https://documentation.bloomreach.com/engagement/reference/export-a-customer-2 https://documentation.bloomreach.com/engagement/reference/export-a-customer-2
16. Does your solution publish transparent, usage-based pricing (events, MTU, storage) in writing, independent of professional-services spend?	No	Bloomreach does not publicly disclose detailed usage-based pricing on its website. Pricing for Bloomreach Engagement is handled via custom quotes and packages rather than a transparent rate card. The website's pricing section encourages contacting Bloomreach for a pricing proposal ("Get Pricing" rather than listing exact costs). While Bloomreach's legal and sales documents define usage metrics like Monthly Processed Events, Monthly Active Users, storage limits, etc., these are used for contracts and overage calculations – the exact prices for those units are not published openly. In summary, you won't find a publicly posted price per event or per user; pricing is tailored and provided through Bloomreach sales, separate from any professional services fees.	https://www.bloomreach.com/en/pricing/engagement https://www.bloomreach.com/en/legal/engagement-appendix

Vendor Deep Dive



Algonomy

Primary Archetype: Commerce & Personalization Engine

Strengths: Retail Specialist; Real-time/Edge

Capability Lens: In-Session Activation

VOX Score: 93.91 / 100

Preference Coverage: 15 / 16 questions “Yes”

📞 Talk to Algonomy: <https://algonomy.com/>

📞 [Ask VOX for Introduction](#)

Oct 2025 CDP Pulse Report - Algonomy

Yes/No	Details	Source(s)	
1. Does your solution provide data-observability dashboards that refresh in ≤ 60 seconds (p95) and show source-to-destination lineage for every first-party data stream?	Yes	Algonomy's Real-time CDP features a live dashboard that provides real-time visibility into data ingestion, event processing, and customer profile counts—aggregated over a 24-hour window. This enables marketing and CX teams to monitor data freshness and system health at a glance. For support teams, the platform provides complete data lineage, allowing them to trace exactly where and why any processing failures occurred. While end users see high-level error messages in real time, the backend diagnostics offer granular, actionable insights for rapid resolution.	https://algonomy.com/customer-data-platform/connectors/
2. Does your solution ship native inbound connectors or APIs for the top-10 commerce, marketing-automation, and service platforms without custom code or PS?	Yes	Algonomy's CDP offers 560+ pre-built connectors across e-commerce, CRM, marketing automation, analytics, etc., allowing quick integration without coding. Users can connect major marketing and commerce apps in minutes via a no-code interface.	https://algonomy.com/customer-data-platform/connectors/
3. Does your solution support server-side and client-side collection for web, mobile-app, email, and in-store (POS/scanner) events in a single SDK?	Yes	Algonomy supports capturing customer events from both online and offline sources in one platform. It ingests data from web and mobile (e.g. site behaviors, app events), email interactions, and offline/POS systems via APIs or SDKs, unifying them in real-time.	https://algonomy.com/customer-data-platform/data-management/
4. Does your solution offer deterministic identity-resolution that stitches online cookies, MAIDs, loyalty IDs and POS receipts into one customer profile?	Yes	Algonomy performs robust identity resolution to create a "Golden Customer Record." It uses deterministic (and probabilistic) matching to merge identifiers across web, mobile, and in-store data into one unified profile. Duplicate records are deduplicated and customer identities (e.g. loyalty IDs, device IDs) are stitched into a single view in real time.	https://algonomy.com/customer-data-platform/ https://algonomy.com/blogs/marketing/why-cdp-enabled-identity-resolution-key-to-individualized-customer-engagement/
5. Does your solution sync online and in-store inventory events in < 5 minutes (p95) to enable BOPIS and endless-aisle use cases?	Yes	Algonomy CDP includes a robust Product Catalog API that supports both inserts and updates, with changes instantly reflected across personalization and campaign modules. It enables real-time synchronization of product data for dynamic content rendering and recommendations.	https://developer.algonomy.com/rcdn/Content/Topics_am_API/Product%20Catalog%20api.htm
6. Does your solution expose decision or personalization APIs that return in < 1 second (p95) so actions fire while the session is active?	Yes	Algonomy's algorithmic decisioning engine is built for real-time, low-latency responses. The platform runs on a globally distributed infrastructure (20+ data centers) to serve personalization results with industry-leading speed. This ensures that recommendation and decision APIs respond almost instantly (well under 1 second in typical scenarios), enabling in-session personalization.	https://algonomy.com/blogs/personalization/ecommerce-spikes-galore-how-is-your-personalization-engine-faring/ https://algonomy.com/blogs/personalization/ecommerce-spikes-galore-how-is-your-personalization-engine-faring/
7. Does your solution include embedded predictive-merchandising models (e.g., next-best product) deployable by business users without data-science help?	Yes	Algonomy provides numerous out-of-the-box AI/ML models tailored for retail (e.g. affinity scoring, replenishment forecasting, product recommendations, churn propensity, CLTV). These predictive models are built into the platform and can be utilized by marketers and merchandisers through the UI, without needing data-scientist intervention. Business users can deploy and leverage "next-best" product recommendations and other predictive insights directly.	https://algonomy.com/customer-data-platform/ https://algonomy.com/algonomy-launches-to-power-digital-first-as-the-new-normal-for-retailers-and-brands-across-the-globe/

Oct 2025 CDP Pulse Report - Algonomy

Yes/No	Details	Source(s)
8. Does your solution allow teams to upload and deploy their own ML/AI models (PMML/ONNX/Python) to score profiles or events in real time?	Yes Algonomy supports PMML-based model imports through its support team. Once uploaded, users can manually trigger or schedule scoring jobs using these models. The system also accommodates external model outputs—enabling marketers or data scientists to ingest scores from models run outside the platform, giving them flexibility in how AI/ML is operationalized.	https://algonomy.com/algonomy-launches-to-power-digital-first-as-the-new-normal-for-retailers-and-brands-across-the-globe/ https://algonomy.com/customer-data-platform/
9. Does your solution provide an SQL workspace or direct lake-house access for segmentation, attribution, and ad-hoc analytics?	Yes Algonomy provides a Customer Lakehouse architecture that gives data teams direct access to customer data for analysis. Users can query the unified data (including raw and processed customer data) via a built-in query engine and connect BI tools (Tableau, PowerBI, etc.) directly to the platform. This enables SQL-like queries for segmentation, attribution, and custom analytics on the CDP data without complex exports.	https://algonomy.com/customer-data-platform/customer-lakehouse/ https://algonomy.com/customer-data-platform/customer-lakehouse/
10. Does your solution run continuous data-quality and anomaly detection with automatic quarantine or alerting before bad data feeds downstream systems?	Yes Automatic quarantining of suspect data is a built-in part of the ingestion pipeline. Users can monitor quarantined records through a dedicated dashboard and download them for remediation. This ensures clean, usable data flows into customer profiles. Additionally, anomaly detection is available out-of-the-box. The platform monitors data at the point of entry, flags irregularities based on HTTP status codes and behavioral thresholds, and raises system alerts—ensuring data fidelity and pipeline stability.	https://algonomy.com/customer-data-platform/customer-lakehouse/ https://algonomy.com/customer-data-platform/connectors/
11. Does your solution enforce role-based access controls and immutable audit logs at the field level across all regions?	Yes Algonomy provides field-level role-based access control (RBAC). Admins can configure access rights down to specific fields by role, user group, or department—ensuring privacy and data governance in large enterprises. The platform also offers an immutable audit trail, accessible via RBAC-defined permissions. This guarantees traceability and compliance with internal audit and regulatory standards.	https://algonomy.com/customer-data-platform/customer-lakehouse/ https://algonomy.com/product-platform/data-infrastructure/
12. Does your solution offer built-in consent & preference management that suppresses activation instantly when permissions change (GDPR/CCPA ready)?	Yes Algonomy's CDP includes integrated consent and privacy controls. It can manage customer consent status and preferences as part of profile data, and will honor those in real time for activation. The platform is designed to comply with GDPR/CCPA – meaning if a user opts out or withdraws consent, the system will immediately stop using that customer's data for targeting/activation.	https://algonomy.com/customer-data-platform/ https://algonomy.com/product-platform/data-infrastructure/
13. Does your solution support regional data residency options (e.g., EU-only, US-only) with encryption at rest (AES-256) and in transit (TLS 1.2+)?	Yes Algonomy offers globally distributed hosting with respect for data residency and sovereignty requirements. They operate many data centers worldwide and can confine data to specific regions to meet local regulations (e.g. EU data stays in EU). All stored customer data is encrypted at rest (using strong encryption), and communications are secured in transit (industry-standard TLS). The platform meets ISO 27001 and similar standards for security.	https://algonomy.com/product-platform/data-infrastructure/ https://algonomy.com/product-platform/data-infrastructure/
14. Does your solution provide reverse-ETL / audience-syndication connectors that push profiles and events to major ad, analytics, and messaging destinations natively?	Yes Algonomy natively supports exporting audiences and events to external systems. With its large library of connectors, users can syndicate customer profiles and segments in real time or batch to advertising platforms, analytics tools, email/SMS services, etc. The platform automates sharing of segments to downstream channels without custom development, functioning effectively as a reverse-ETL for activation.	https://algonomy.com/customer-data-platform/ https://algonomy.com/customer-data-platform/data-management/

Oct 2025 CDP Pulse Report - Algonomy

Yes/No	Details	Source(s)		
15. Does your solution expose a unified customer-profile API/stream that any downstream system can query or subscribe to in real time?	Yes	<p>Algonomy exposes unified customer profiles and segments through real-time APIs and streams. Downstream systems can query the CDP (or subscribe to updates) to obtain the latest customer attributes or segment memberships on the fly. In practice, Algonomy's Audience Manager streams profile segments in real time so external tools always have up-to-date data.</p>	https://algonomy.com/customer-data-platform/data-management/	https://www.g2.com/products/algonomy-omni-channel-personalization-suite/reviews
16. Does your solution publish transparent, usage-based pricing (events, MTU, storage) in writing, independent of professional-services spend?	No	<p>Algonomy does not publish pricing details on its website. Pricing is not disclosed as usage-based rates; it typically requires contacting the vendor. Pricing model is based on three key dimensions: profile counts, event volumes and selected enterprise features/modules.</p> <p>Algonomy offers flexible pricing tailored to each customer's data scale and functional requirements, with transparent tiering that aligns to business growth.</p>	https://www.getapp.com/customer-management-software/a/algonomy-customer-data-platform/	

Customer Data Infrastructure / Composable



Vendor Deep Dive



RudderStack

Primary Archetype: Customer Data Infrastructure / Composable

Strengths: Warehouse-native/Composable;
Real-time/Edge

Capability Lens: Data Visibility & Quality; Security,
Consent & Residency; Customization
& Composability

VOX Score: 93.98 / 100

Preference Coverage: 15 / 16 questions “Yes”

📞 Talk to RudderStack: <https://www.rudderstack.com/>

📞 [Ask VOX for Introduction](#)

Oct 2025 CDP Pulse Report - RudderStack

Yes/No	Details	Source(s)
1. Does your solution provide data-observability dashboards that refresh in ≤ 60 seconds (p95) and show source-to-destination lineage for every first-party data stream?	Yes RudderStack offers real-time observability for event pipelines. It provides dashboards (e.g. a Grafana-based Health Dashboard) that update live with metrics on events ingested and delivered. Users get end-to-end visibility into data flows – including a lineage view of events from sources through to destinations – to monitor every first-party data stream in near real time.	rudderstack.com/blog/real-time-data-integration/ rudderstack.com/use-case/data-quality/
2. Does your solution ship native inbound connectors or APIs for the top-10 commerce, marketing-automation, and service platforms without custom code or PS?	Yes RudderStack has an extensive integration library (200+ connectors) covering all major commerce, marketing, and service platforms out-of-the-box. For example, it provides native connectors for e-commerce platforms like Shopify, marketing automation tools like Marketo and Mailchimp, CRM and customer service systems like Salesforce, HubSpot, Zendesk, etc. These sources/destinations can be set up without any custom coding or professional services.	rudderstack.com/integration/ rudderstack.com/integration/
3. Does your solution support server-side and client-side collection for web, mobile-app, email, and in-store (POS/scanner) events in a single SDK?	Yes RudderStack supports both client-side and server-side data collection across channels. It provides open-source SDKs for websites (JavaScript), mobile apps (iOS, Android, React Native, etc.), and server applications (Node.js, Python, etc.), allowing unified event tracking. This means events from web or mobile frontends, email or backend systems, and even in-store/POS data (via server APIs or SDKs) can all be captured and sent through one platform.	https://www.rudderstack.com/docs/sources/event-streams/sdks/ https://www.rudderstack.com/docs/sources/event-streams/sdks/
4. Does your solution offer deterministic identity-resolution that stitches online cookies, MAIDs, loyalty IDs and POS receipts into one customer profile?	Yes RudderStack's Profiles feature provides deterministic identity resolution ("identity stitching"). It builds an identity graph in your data warehouse that can merge identifiers from multiple sources into a single customer profile. In practice, this means it can take online identifiers (cookies or mobile Ad IDs) together with offline IDs (loyalty program IDs, POS transaction IDs, etc.) and reliably link them to the same user record, yielding a unified customer 360 view.	rudderstack.com/use-case/identity-resolution/ rudderstack.com/use-case/identity-resolution/
5. Does your solution sync online and in-store inventory events in < 5 minutes (p95) to enable BOPIS and endless-aisle use cases?	Yes RudderStack's event streaming pipeline operates in real time, typically processing and delivering events within seconds. It can ingest inventory updates from online systems and in-store/point-of-sale systems and sync them to downstream tools very quickly (well under a 5-minute 95th percentile). This low latency data flow supports retail use cases like "Buy Online, Pick Up In Store" and endless aisle, ensuring inventory changes are reflected across channels almost immediately.	rudderstack.com/blog/real-time-data-integration/ rudderstack.com/blog/real-time-data-integration/
6. Does your solution expose decision or personalization APIs that return in <1 second (p95) so actions fire while the session is active?	Yes RudderStack provides a high-performance Profiles Activation API for real-time personalization queries. This API lets downstream applications retrieve a user's profile traits or computed features from a fast in-memory store (Redis) with sub-second latency. In practice, it enables session-level decisions (e.g. personalized content or offers) to be made instantly while the user is still active on the site or app.	https://www.rudderstack.com/docs/profiles/activation-api/ https://www.rudderstack.com/docs/profiles/activation-api/
7. Does your solution include embedded predictive-merchandising models (e.g., next-best product) deployable by business users without data-science help?	Yes RudderStack includes Predictions (Propensity Scores) as part of its platform – a set of built-in ML models that non-technical users can leverage. Currently, it offers pre-packaged models like churn propensity and lead conversion scoring. These models are warehouse-native and can be configured and deployed via a user-friendly UI by analysts or marketers without writing code or requiring a dedicated data science pipeline.	rudderstack.com/blog/early-access-feature-launch-rudderstack-predictions/ rudderstack.com/blog/early-access-feature-launch-rudderstack-predictions/

Oct 2025 CDP Pulse Report - RudderStack

Yes/No	Details	Source(s)
8. Does your solution allow teams to upload and deploy their own ML/AI models (PMML/ONNX/Python) to score profiles or events in real time?	Yes RudderStack's platform is extensible to custom models. Data teams can develop their own machine learning models in Python as part of the warehouse-centric Profiles workflow. The Predictions feature lets you train and deploy these custom models (for example, predicting any custom outcome like LTV or propensity) on your data. The model scoring runs in real time on your warehouse data, and RudderStack monitors model performance (with automated retraining for drift) – all within your infrastructure.	rudderstack.com/product/predictions/ rudderstack.com/blog/early-access-feature-launch-rudderstack-predictions/
9. Does your solution provide an SQL workspace or direct lake-house access for segmentation, attribution, and ad-hoc analytics?	Yes RudderStack is built around your data warehouse/lake (“warehouse-first”), so all customer data is available directly in databases like Snowflake, BigQuery, etc. This means teams can use SQL on their own warehouse for segmentation, attribution, and any ad-hoc analysis. Moreover, RudderStack offers tools like Audiences (a UI to create segments using warehouse data without SQL) and allows defining custom models or traits via SQL in its Profiles/Reverse ETL module, giving flexibility to analysts and engineers alike.	rudderstack.com/use-case/identity-resolution/ https://www.rudderstack.com/docs/data-pipelines/reverse-etl/features/audiences/
10. Does your solution run continuous data-quality and anomaly detection with automatic quarantine or alerting before bad data feeds downstream systems?	Yes RudderStack includes an automated Data Quality toolkit as part of its governance features. Users can define Tracking Plan rules (expected event schemas and properties), and RudderStack will continuously monitor incoming events for violations. If an event doesn't conform (bad data or anomaly), it can automatically quarantine or drop that event, or reroute it, before it reaches any destination. The system also provides alerts and a health dashboard so data teams are notified of schema drift or anomalies in real time, preventing bad data from propagating.	rudderstack.com/use-case/data-quality/ rudderstack.com/use-case/data-quality/
11. Does your solution enforce role-based access controls and immutable audit logs at the field level across all regions?	Yes RudderStack provides robust security and access controls. It has organization and workspace-level role-based access control (with roles like Admin, Member, and more granular resource-level permissions) that can restrict who can view or edit specific data pipelines or even sensitive fields (PII) in event payloads. All changes and user actions in the system are recorded in detailed audit logs that cannot be altered, ensuring an immutable trail for compliance. These security measures are in effect for all deployments (including multi-region environments) and help meet requirements for standards like SOC 2, GDPR, and HIPAA.	https://www.rudderstack.com/docs/dashboard-guides/permissions-management/ https://www.rudderstack.com/docs/dashboard-guides/permissions-management/
12. Does your solution offer built-in consent & preference management that suppresses activation instantly when permissions change (GDPR/CCPA ready)?	No RudderStack does not have a built-in end-user consent preference center. Instead, it integrates with existing consent management platforms and mechanisms. You can pass consent flags from popular CMPs (like OneTrust, Ketch, etc.) into RudderStack's SDK, and RudderStack will honor those flags by blocking or filtering events to destinations as needed. It also provides a User Suppression API to honor “do not track” or deletion requests. Changes in consent are propagated – for example, if a user opts out, RudderStack can stop forwarding that user's data – but the solution relies on external consent management or custom implementation to capture the preference. In summary, it is GDPR/CCPA compliant when properly integrated, but it doesn't natively manage the consent UI or workflow; it uses the consent info you supply to instantly suppress or delete data.	https://www.rudderstack.com/docs/data-governance/consent-management/

Oct 2025 CDP Pulse Report - RudderStack

	Yes/No	Details	Source(s)
13. Does your solution support regional data residency options (e.g., EU-only, US-only) with encryption at rest (AES-256) and in transit (TLS 1.2+)?	Yes	RudderStack supports deploying the platform in different regions to satisfy data residency requirements. For example, customers can use a European data plane domain to keep data in the EU, or a US data plane for US-only data, etc. All data handled by RudderStack is secured via strong encryption – data in transit is sent exclusively over HTTPS using TLS 1.2+, and data at rest in RudderStack's managed services (datastores, backups) is encrypted using AES-256.	https://funnelfox.com/docs/integrations/analytics/rudderstack_rudderstack.com/data-transfer-impact-assessment/
14. Does your solution provide reverse-ETL / audience-syndication connectors that push profiles and events to major ad, analytics, and messaging destinations natively?	Yes	Yes. RudderStack includes a native Reverse ETL capability (warehouse-to-any integration) that can sync computed customer data out to hundreds of destinations. This means you can take unified profiles or audience segments from your warehouse and push them to advertising platforms (Google, Facebook, TikTok, etc.), analytics tools, email marketing services, CRMs, and other downstream applications. The platform provides many built-in connectors for these “audience syndication” use cases – including support for transmitting audiences to ad networks or customer lists to CRM and messaging tools – without custom development.	https://www.rudderstack.com/docs/data-pipelines/reverse-etl/features/audiences/ https://www.rudderstack.com/docs/data-pipelines/reverse-etl/features/audiences/
15. Does your solution expose a unified customer-profile API/stream that any downstream system can query or subscribe to in real time?	Yes	RudderStack exposes unified customer profiles through its Profiles API endpoints. In practice, after building a 360° customer profile (via identity resolution and computed traits in the warehouse), you can query that profile's attributes in real time using RudderStack's Activation API. This API allows other services to pull down the latest profile data (e.g. for personalization or customer service) on demand. Additionally, RudderStack can forward real-time event streams, so systems can subscribe to live event data or profile updates as they flow through (via webhooks or streaming destinations), ensuring up-to-date information in all tools.	https://www.rudderstack.com/docs/profiles/dev-docs/activation-api/ https://www.rudderstack.com/docs/profiles/dev-docs/activation-api/
16. Does your solution publish transparent, usage-based pricing (events, MTU, storage) in writing, independent of professional-services spend?	Yes	RudderStack's pricing is openly published and based on product usage (primarily event volume tiers). The company offers a Free plan (e.g. up to 250k events/month free) and clearly defined paid plans that scale by monthly event count (Starter, Growth, etc.), with the option to customize for higher volumes. These prices and limits are provided on the website, and are not tied to professional services commitments. In other words, you pay for the data volume and features you use; professional services are optional and separate.	rudderstack.com/pricing/ rudderstack.com/pricing/

Vendor Deep Dive



Tealium

Primary Archetype: Customer Data Infrastructure / Composable

Strengths: Real-time/Edge; Governance & Residency

Capability Lens: Data Visibility & Quality; Security, Consent & Residency; Customization & Composability

VOX Score: 81.06 / 100

Preference Coverage: 13 / 16 questions “Yes”

📞 Talk to Tealium: <https://tealium.com/>

📞 [Ask VOX for Introduction](#)

Oct 2025 CDP Pulse Report - Tealium

Yes/No	Details	Source(s)
1. Does your solution provide data-observability dashboards that refresh in ≤ 60 seconds (p95) and show source-to-destination lineage for every first-party data stream?	Yes Tealium offers built-in operational dashboards (e.g. the Data Supply Chain and Connectors dashboards) that give real-time visibility into data flows. Users can monitor each data source through to each destination with health metrics and lineage. These dashboards update in near real-time, enabling quick detection of issues in first-party data streams.	https://tealium.com/products/integrations/ https://tealium.com/products/integrations/
2. Does your solution ship native inbound connectors or APIs for the top-10 commerce, marketing-automation, and service platforms without custom code or PS?	Yes Tealium provides an extensive Integrations Marketplace with over 1,300 pre-built connectors covering all major marketing, commerce, analytics, advertising, CRM, and service platforms. These native integrations require no custom coding or professional services to set up, allowing data to flow from popular tools (e.g. e-commerce, email marketing, CRM) directly into Tealium and vice versa.	https://www.uniphore.com/products/marketing-ai/cdp-agent/integrations/ https://www.uniphore.com/products/marketing-ai/cdp-agent/integrations/
3. Does your solution support server-side and client-side collection for web, mobile-app, email, and in-store (POS/scanner) events in a single SDK?	Yes Tealium supports unified data collection across web, mobile apps, email, and even offline POS or other devices. It offers both client-side tag/SDK integrations and server-side APIs so that all these event sources feed into one centralized Customer Data Hub. This means a consistent tracking framework for online and offline channels. For example, Tealium can ingest web browser events, in-app mobile events, and offline transactions into the same system in real time.	https://tealium.com/products/integrations/
4. Does your solution offer deterministic identity-resolution that stitches online cookies, MAIDs, loyalty IDs and POS receipts into one customer profile?	Yes Tealium AudienceStream CDP provides deterministic identity resolution via its patented Visitor Stitching technology. It can unify identifiers like web cookies, mobile device IDs (MAIDs), loyalty or account IDs, and offline transaction IDs into a single profile. As soon as a common identifier is observed (e.g. a user logs in or provides an email/loyalty ID), Tealium stitches the previously separate profiles together in real time to form one unified customer profile.	https://tealium.com/resource/datasheet/visitor-stitching/ https://tealium.com/resource/datasheet/visitor-stitching/
5. Does your solution sync online and in-store inventory events in < 5 minutes (p95) to enable BOPIS and endless-aisle use cases?	Yes Tealium can ingest and unify both online and offline (in-store) events in real time. For example, point-of-sale transactions or inventory updates from stores can be fed into Tealium via APIs or file feeds as they occur. Once the data hits Tealium's CDP (typically within seconds), it becomes instantly available to update customer profiles or trigger actions. This rapid synchronization enables use cases like BOPIS and endless aisle, where inventory changes are reflected across channels almost immediately.	https://tealium.com/download/tealium-real-time-for-real-action/
6. Does your solution expose decision or personalization APIs that return in < 1 second (p95) so actions fire while the session is active?	Yes Tealium provides high-performance APIs for real-time personalization and deciding. Notably, the Tealium Moments API allows on-demand retrieval of unified profile attributes in extremely low latency (tens of milliseconds). This API can be called from a webpage or app during a live session to decide, for example, which experience to show, and it responds well under 1 second (in fact, ~0.06s on average for typical payloads), ensuring in-session actions can fire immediately.	https://tealium.com/products/tealium-moments/moments-api/ https://tealium.com/products/tealium-moments/moments-api/
7. Does your solution include embedded predictive-merchandising models (e.g., next-best product) deployable by business users without data-science help?	Yes Tealium offers an embedded machine learning feature called Predict ML as part of its CDP. This allows non-technical users to create and deploy custom predictive models (such as likelihood to purchase or churn) with just a few clicks, no coding needed. The UI guides marketers to train models using their customer data and then automatically generates propensity scores that can drive "next-best" product recommendations or other campaigns – all without requiring a data scientist.	https://tealium.com/resource/datasheet/tealium-predict-ml/

Oct 2025 CDP Pulse Report - Tealium

Yes/No	Details	Source(s)	
8. Does your solution allow teams to upload and deploy their own ML/AI models (PMML/ONNX/Python) to score profiles or events in real time?	No	Tealium's approach to machine learning is to provide built-in Predict ML capabilities rather than enabling custom model uploads. It does not natively support importing external ML models in formats like PMML or ONNX for deployment within the platform. Instead, Tealium focuses on its integrated ML, where models are trained on data inside the CDP and then applied to profiles/events. Organizations needing custom algorithms would typically score data outside Tealium and send the results into the CDP.	https://customerthink.com/new-tealium-predict-machine-learning-solution-helps-marketers-proactively-understand-cx-data-to-deliver-roi-faster/
9. Does your solution provide an SQL workspace or direct lake-house access for segmentation, attribution, and ad-hoc analytics?	Yes	Tealium provides DataAccess features that give direct access to the underlying event and profile data. For example, Tealium can automatically warehouse data in systems like Amazon Redshift (AudienceDB/EventDB) or Snowflake, where clients can connect with SQL tools. This means analysts can run SQL queries directly on the raw Tealium data (or even connect BI tools) to build custom segments, perform multi-touch attribution analysis, and conduct ad-hoc queries beyond the CDP's UI.	https://docs.tealium.com/server-side/data-storage/audiencedb-eventdb/about/
10. Does your solution run continuous data-quality and anomaly detection with automatic quarantine or alerting before bad data feeds downstream systems?	No	Tealium includes tools for data quality monitoring (such as Event Specifications that validate incoming event payloads in real time and flag schema mismatches). However, it does not automatically quarantine or block "bad" data nor perform AI-driven anomaly detection out-of-the-box. Invalid events are flagged in dashboards for review but still pass through the system. Proactive anomaly detection/alerting would require additional configuration or external tools.	https://docs.tealium.com/server-side/event-specifications/about/
11. Does your solution enforce role-based access controls and immutable audit logs at the field level across all regions?	Yes	Tealium provides robust role-based access control (RBAC) with fine-grained permissions. Administrators can assign pre-defined roles or custom permission sets to users, controlling access down to specific profiles and even parts of the data layer. Tealium iQ, for example, supports locking specific data layer elements or tags so only authorized roles can modify themtealium.com. All changes and publish actions are tracked in version histories, providing an audit trail of who changed what.	https://tealium.com/resource/datasheet/tealium-data-security-and-privacy-tools-controls-and-infrastructure/ https://tealium.com/resource/datasheet/tealium-data-security-and-privacy-tools-controls-and-infrastructure/
12. Does your solution offer built-in consent & preference management that suppresses activation instantly when permissions change (GDPR/CCPA ready)?	Yes	Yes. Tealium includes built-in Consent Management capabilities. It can capture users' consent preferences and enforce them in real time. For example, if a user opts out, Tealium's system will immediately stop triggering any connector actions or tags categorized under the opt-out purposes. Tealium supports GDPR/CCPA requirements by allowing consent categories, and each integration action is tied to consent categories so that when consent is withdrawn, data activation is suppressed instantly.	https://docs.tealium.com/server-side/connectors/about/

Oct 2025 CDP Pulse Report - Tealium

Yes/No	Details	Source(s)	
13. Does your solution support regional data residency options (e.g., EU-only, US-only) with encryption at rest (AES-256) and in transit (TLS 1.2+)?	Yes	Tealium offers global data centers and lets customers choose their data hosting region to meet data residency requirements (for instance, keeping data in the EU only or US only). All customer data is protected with strong encryption – both at rest in Tealium's databases and in transit over networks. Tealium adheres to industry-standard ciphers (AES-256 for data at rest and TLS 1.2+ for data in motion), ensuring compliance with security and privacy regulations in different regions.	https://tealium.com/protecting-your-greatest-competitive-advantage-data/ https://tealium.com/protecting-your-greatest-competitive-advantage-data/
14. Does your solution provide reverse-ETL / audience-syndication connectors that push profiles and events to major ad, analytics, and messaging destinations natively?	Yes	Tealium's 1,300+ integration connectors cover essentially all major downstream systems for audiences and events. The platform can natively push enriched profiles or event data to advertising networks (Google, Facebook, etc.), analytics platforms, email/SMS marketing tools, and other martech systems in real time. These audience syndication connectors are configured without coding – users simply map profile attributes or segments to the target platform, and Tealium handles the delivery.	https://tealium.com/products-integrations/
15. Does your solution expose a unified customer-profile API/stream that any downstream system can query or subscribe to in real time?	Yes	Yes. Tealium provides a unified profile API (e.g. the Tealium Moments API, part of its Customer Data Hub). This API allows external systems to query a customer's up-to-date profile data on demand (e.g., retrieve segment memberships or attributes as JSON) with very low latencytealium.com. Additionally, Tealium's event feeds and connectors enable pushing a real-time stream of profile changes or events out to other systems (or data warehouses) if a subscribe model is preferred.	https://tealium.com/products/tealium-moments/moments-api/ https://tealium.com/products/tealium-moments/moments-api/
16. Does your solution publish transparent, usage-based pricing (events, MTU, storage) in writing, independent of professional-services spend?	No	Tealium's pricing is not published openly on its website. Pricing is typically customized per customer and based on usage factors like event volumes or audience size, often negotiated via the sales team. In practice, Tealium's license is usage-based (for example, a tier of event counts includes certain service support hoursg2.com), but there is no public, self-serve pricing page or fixed rate card independent of deals that may include software and service bundles.	https://www.g2.com/compare/tealium-customer-data-hub-vs-twilio-segment

Vendor Deep Dive



Netcore

Primary Archetype: Customer Data Infrastructure / Composable

Strengths: Warehouse-native/Composable;
Real-time/Edge

Capability Lens: Security, Consent & Residency;
Customization & Composability

VOX Score: 62.32 / 100

Preference Coverage: 10 / 16 questions “Yes”

📞 Talk to Netcore: <https://netcorecloud.com/>

📞 [Ask VOX for Introduction](#)

Oct 2025 CDP Pulse Report - Netcore

	Yes/No	Details	Source(s)
1. Does your solution provide data-observability dashboards that refresh in ≤ 60 seconds (p95) and show source-to-destination lineage for every first-party data stream?	No	Netcore's CDP offers real-time dashboards and live insights (e.g. active users, conversion events) for marketing performance. However, it does not explicitly document sub-60s refresh rates or end-to-end data lineage visualization for each data stream. The focus is on marketing KPIs and user behavior dashboards, not low-level pipeline observability.	https://netcorecloud.com/customer-engagement/dashboards-insights/ https://netcorecloud.com/cdp/analytics-insights/
2. Does your solution ship native inbound connectors or APIs for the top-10 commerce, marketing-automation, and service platforms without custom code or PS?	Yes	Netcore provides a broad set of native connectors and APIs for popular systems. For example, it integrates with e-commerce platforms like Shopify and Magento, CRM/marketing systems like Salesforce Commerce, mobile attribution (Appsflyer), databases, cloud storage, etc., all without custom coding. These built-in connectors enable ingesting data from top commerce and marketing platforms (and even custom sources via API) out of the box.	https://netcorecloud.com/blog/clevertap-alternatives/ https://netcorecloud.com/customer-data-platform
3. Does your solution support server-side and client-side collection for web, mobile-app, email, and in-store (POS/scanner) events in a single SDK?	Yes	Netcore's "full-stack" CDP supports data collection from client-side sources (web JavaScript SDK, iOS/Android app SDKs) as well as server-side uploads. It lists connectors for websites, mobile apps, and even offline sources (e.g. SFTP file or API imports for POS data). All these feed into one unified platform. (Email engagement is also tracked through Netcore's email product, ensuring events from email are captured in the same customer data platform.)	https://netcorecloud.com/customer-data-platform https://netcorecloud.com/customer-data-platform
4. Does your solution offer deterministic identity-resolution that stitches online cookies, MAIDs, loyalty IDs and POS receipts into one customer profile?	Yes	Netcore Cloud includes an Identity Resolution engine that unifies identifiers across devices and touchpoints into a single profile. It uses a proprietary ID graph algorithm to stitch customer behavior from cookies, mobile IDs, offline identifiers, etc., under one hood. This deterministic identity resolution runs in real time, breaking down data silos and giving a 360° view of each customer.	https://netcorecloud.com/cdp/identity-resolution/ https://netcorecloud.com/cdp/identity-resolution/
5. Does your solution sync online and in-store inventory events in < 5 minutes (p95) to enable BOPIS and endless-aisle use cases?	No	There is no clear documentation of a guaranteed sub-5-minute inventory sync SLA. Netcore can ingest and act on inventory updates (it even supports automated triggers like "back-in-stock" alerts for when products return to inventory, which suggests near-real-time handling). However, the vendor does not explicitly state a p95 <5min sync latency for online vs. in-store inventory events. BOPIS/endless-aisle scenarios are possible, but no specific timing commitment is published.	https://netcorecloud.com/plan-and-pricing/
6. Does your solution expose decision or personalization APIs that return in <1 second (p95) so actions fire while the session is active?	Yes	Netcore offers real-time personalization capabilities that imply very fast response times. For instance, its product includes AI-powered site search and real-time product recommendations which are delivered on-the-fly during user sessions. While exact p95 latency isn't quoted, the platform is built for in-session decisions – e.g. an API powering personalization or recommendation widgets would respond virtually instantly (within a second) to personalize content while the visitor is active.	https://netcorecloud.com/infrastructure-security/
7. Does your solution include embedded predictive-merchandising models (e.g., next-best product) deployable by business users without data-science help?	Yes	Netcore's CDP comes with built-in AI/ML models for merchandising and customer analytics that non-technical users can leverage. It provides out-of-the-box product recommendations ("next best product") driven by custom RFM and AI models, as well as predictive segments and churn propensity models. Marketers can simply input parameters (e.g. product SKUs) to get segment or recommendation outputs – no data science coding needed. These embedded models (for churn, CLTV, best channel/time to send, etc.) are part of the platform's AI toolkit.	https://netcorecloud.com/customer-data-platform https://netcorecloud.com/infrastructure-security/

Oct 2025 CDP Pulse Report - Netcore

Yes/No	Details	Source(s)	
8. Does your solution allow teams to upload and deploy their own ML/AI models (PMML/ONNX/Python) to score profiles or events in real time?	No	Netcore's platform does not advertise support for importing custom ML models. Its approach is to provide built-in AI models (for personalization, predictions, etc.) rather than letting users deploy their own PMML/ONNX/Python models. We found no documentation of any "bring your own model" feature – users are expected to use the native AI capabilities (e.g. Netcore's proprietary algorithms for scoring and predictions) rather than plug in external models.	https://netcorecloud.com/customer-data-platform
9. Does your solution provide an SQL workspace or direct lake-house access for segmentation, attribution, and ad-hoc analytics?	No	Netcore does not offer a user-facing SQL query workspace or direct access to a cloud data lake. Instead, it emphasizes visual analytics dashboards and built-in segmentation tools. The platform can export event data (an "Event Data Export" feature is available as an add-on) for external analysis, but there is no indication of a live SQL IDE or lake-house connectivity for ad-hoc queries. Users rely on Netcore's own reporting UI and can request data exports if needed.	https://netcorecloud.com/plan-and-pricing/
10. Does your solution run continuous data-quality and anomaly detection with automatic quarantine or alerting before bad data feeds downstream systems?	No	Netcore Cloud does not list a dedicated data quality monitoring module that quarantines bad data. It does have anomaly detection as part of its AI insights ("Raman") – but that is geared toward spotting anomalies in marketing KPIs/time-series performance (e.g. unusual drops or spikes in engagement), not schema or data integrity errors in real-time feeds. There's no mention of automatically quarantining bad inbound data or halting downstream feeds; the focus is on alerting marketers to performance anomalies rather than enforcing data-quality gates.	https://netcorecloud.com/blog/ramans-intelligent-insights-understanding-the-machine-learning-models-that-enable-anomaly-detection/
11. Does your solution enforce role-based access controls and immutable audit logs at the field level across all regions?	Yes	Netcore implements granular role-based access control (RBAC) for data and features. Field-level access restrictions and admin controls are supported (the platform includes Role-Based Access Management as a core feature). While the documentation doesn't explicitly mention "immutable field-level audit logs," Netcore's security practices include extensive logging/monitoring of access and changes to protect data integrity. It is also compliant with industry standards (ISO 27001, etc.), implying that administrative actions are auditable even across its global deployments.	https://netcorecloud.com/plan-and-pricing/ https://netcorecloud.com/security/
12. Does your solution offer built-in consent & preference management that suppresses activation instantly when permissions change (GDPR/CCPA ready)?	Yes	Yes – Netcore Cloud's platform is GDPR-ready and provides features for consent and preference management. It enables its clients to honor user consent changes in real-time. For example, the system can instantly suppress or halt messaging to a user who withdraws consent. Netcore highlights that its platform has capabilities to help customers remain compliant with GDPR (and similar laws) by managing data subject preferences and rights easily. (The product includes GDPR compliance tools out-of-the-box, so preference changes are respected immediately in all activation channels.)	https://netcorecloud.com/security/ https://slashdot.org/software/comparison/Adjust-vs-Netcore/
13. Does your solution support regional data residency options (e.g., EU-only, US-only) with encryption at rest (AES-256) and in transit (TLS 1.2+)?	Yes	Netcore Cloud can accommodate data residency requirements and adheres to strong encryption standards. They host infrastructure across multiple regions/data centers (which are ISO 27001 and SOC2 certified) to meet local data storage needs. All data at rest is encrypted (using modern protocols analogous to AES-256) and data in transit is secured via TLS 1.2+. (The security documentation explicitly notes TLS 1.2/1.3 for data in motion and advanced encryption for data at rest.)	https://netcorecloud.com/security/ https://netcorecloud.com/security/
14. Does your solution provide reverse-ETL / audience-syndication connectors that push profiles and events to major ad, analytics, and messaging destinations natively?	Yes	Yes. Netcore natively supports audience syndication to external platforms. The CDP can push customer data (segments, events) to popular ad and marketing channels. For instance, it integrates with Facebook/Instagram Custom Audiences and Google Ads Customer Match for audience uploads. It also connects with analytics tools (e.g. Google Analytics) and messaging channels (WhatsApp, email, SMS through its own modules) directly. These out-of-the-box connectors let you activate segments in third-party adtech/martech systems without needing custom middleware.	https://netcorecloud.com/plan-and-pricing/ https://netcorecloud.com/customer-data-platform

Oct 2025 CDP Pulse Report - Netcore

Yes/No	Details	Source(s)	
15. Does your solution expose a unified customer-profile API/stream that any downstream system can query or subscribe to in real time?	Yes	Netcore's CDP creates a unified profile for each customer and makes the consolidated data accessible. Clients (controllers) can retrieve profile data (identifiers, attributes, event history) on demand – for example, Netcore notes it will export a data subject's full profile information upon request. This implies an API or data service to query profiles. While not branded as a “profile stream,” the ability to pull a unified customer profile via API is supported, enabling downstream systems to query the latest profile data in real time.	https://netcorecloud.com/gdpr/ https://netcorecloud.com/security/
16. Does your solution publish transparent, usage-based pricing (events, MTU, storage) in writing, independent of professional-services spend?	No	Netcore does not publicly disclose detailed usage-based pricing on its website. Its pricing page offers plan tiers (Startup, Growth, Enterprise) and emphasizes custom solutions rather than listing per-event or per-user costs. There is no transparent rate card for events, monthly tracked users, or data storage. Pricing appears to be provided via sales consultation or predefined bundles (and is not tied to professional services spend, which is separate). In summary, the exact usage-based pricing metrics are not published in writing.	https://netcorecloud.com/plan-and-pricing/ https://netcorecloud.com/plan-and-pricing/

Vendor Deep Dive



Treasure Data

Primary Archetype: Customer Data Infrastructure / Composable

Strengths: Warehouse-native/Composable;
Governance & Residency

Capability Lens: Data Visibility & Quality; Security,
Consent & Residency; Customization
& Composability

VOX Score: 93.91 / 100

Preference Coverage: 15 / 16 questions “Yes”

📞 Talk to Treasure Data: <https://www.treasuredata.com/>

📞 [Ask VOX for Introduction](#)

Oct 2025 CDP Pulse Report - Treasure Data

	Yes/No	Details	Source(s)
1. Does your solution provide data-observability dashboards that refresh in ≤ 60 seconds (p95) and show source-to-destination lineage for every first-party data stream?	Yes	Treasure Data's CDP includes an observability layer for full transparency into data pipelines. Users can monitor data ingestion and delivery in near real time, seeing how first-party data flows from sources into the CDP and out to destinations. The platform supports true real-time streaming data collection and provides dashboards/logs to troubleshoot and trace data movements.	https://www.treasuredata.com/technology/
2. Does your solution ship native inbound connectors or APIs for the top-10 commerce, marketing-automation, and service platforms without custom code or PS?	Yes	Treasure Data offers 400+ pre-built data connectors covering popular commerce, marketing automation, CRM, and service platforms. These native inbound integrations (and APIs) allow code-free ingestion from leading systems (e.g. e-commerce, email, ads, CRM tools) into the CDP. No custom coding or professional services are required to connect most major third-party data sources.	https://www.treasuredata.com/technology/ https://www.treasuredata.com/partnerships
3. Does your solution support server-side and client-side collection for web, mobile-app, email, and in-store (POS/scanner) events in a single SDK?	Yes	Treasure Data supports unified data collection across channels. It provides client-side SDKs for websites (JavaScript) and mobile apps (iOS, Android, Unity, etc.) as well as server-side and batch ingestion methods. This means online events (web/mobile) and offline events (POS transactions, loyalty scans, etc.) can all be captured into one platform. The CDP accepts offline or online, batch or real-time data, with SDKs and connectors to handle everything from web clicks to in-store sales.	https://treasure-data.github.io/td-nibs/architecture.html
4. Does your solution offer deterministic identity-resolution that stitches online cookies, MAIDs, loyalty IDs and POS receipts into one customer profile?	Yes	Treasure Data's CDP performs deterministic identity resolution (with optional probabilistic methods) to build unified customer profiles. It can stitch together identifiers such as web cookies, mobile device IDs, email/loyalty IDs, and offline purchase records into a single "golden" customer profile. The platform assigns persistent customer IDs and uses robust matching algorithms to ensure that online and in-store identifiers for the same person are merged into one profile.	https://www.treasuredata.com/blog/the-marketers-guide-to-identity-resolution-is-your-data-hiding-things-from-you/
5. Does your solution sync online and in-store inventory events in < 5 minutes (p95) to enable BOPIS and endless-aisle use cases?	Yes	Treasure Data can handle fast integration of inventory data from both online and in-store systems. Its streaming ingestion framework allows near real-time updates (on the order of minutes or faster) to ensure inventory changes are quickly reflected. This enables retail use cases like "Buy Online, Pick Up In Store" and endless aisle, where the CDP can unify and update product availability across channels almost instantly to power up-to-date customer experiences.	https://www.treasuredata.com/technology/ https://www.treasuredata.com/blog/from-bopus-to-mopus-what-the-rise-of-in-store-pickup-means-for-retail/
6. Does your solution expose decision or personalization APIs that return in <1 second (p95) so actions fire while the session is active?	Yes	Treasure Data exposes real-time profile query APIs to drive personalization and decisions within active user sessions. For example, its Profile API allows external systems (e.g. a website or app) to request a customer's unified profile or segment info on the fly and get a sub-second response. This enables on-site personalization or decisioning (such as next-best offer) to occur immediately while the customer is still engaged, meeting low-latency requirements.	
7. Does your solution include embedded predictive-merchandising models (e.g., next-best product) deployable by business users without data-science help?	Yes	Treasure Data's CDP includes built-in AI/ML capabilities for marketers, such as predictive scoring and recommendation models. Non-technical users can leverage these to identify things like next-best product or churn risk without coding. The system can automatically analyze customer attributes/behavior and generate predictions (e.g. product affinity or likelihood to buy) that business users can deploy in campaigns. This "auto-ML" approach means users can configure and use predictive merchandising models via a UI, without needing a data scientist to build models from scratch.	https://www.treasuredata.com/blog/personalization-with-a-cdp/ https://www.treasuredata.com/blog/the-marketers-guide-to-identity-resolution-is-your-data-hiding-things-from-you/

Oct 2025 CDP Pulse Report - Treasure Data

Yes/No	Details	Source(s)
8. Does your solution allow teams to upload and deploy their own ML/AI models (PMML/ONNX/Python) to score profiles or events in real time?	Yes Treasure Data supports a “bring your own model” approach for advanced teams. Data scientists can develop custom machine learning models (e.g. in Python or exported in standard formats) and run or integrate them within the CDP environment. The platform’s AI “Foundry” and workflow tools let you deploy your own prediction functions to score customers or events using custom code, all operating on Treasure Data’s infrastructure. This flexibility means teams can use their preferred ML libraries or pre-trained models and have Treasure Data execute the scoring in real time on incoming data.	https://www.treasuredata.com/solutions/data-science/
9. Does your solution provide an SQL workspace or direct lake-house access for segmentation, attribution, and ad-hoc analytics?	Yes Treasure Data’s platform includes a full SQL query engine and workspace for analysts. Users can run ad-hoc queries (using Presto/Trino SQL) directly on the unified customer dataset to create segments, perform multi-touch attribution, or do custom analytics. The CDP’s underlying cloud data lake is queryable in-place (even supporting a “zero copy” approach to query external data), so teams have direct analytical access without needing to export data. This SQL workspace enables flexible, in-depth analysis on all customer data within the CDP.	https://docs.treasuredata.com/articles/pd/querying-data https://www.treasuredata.com/technology/
10. Does your solution run continuous data-quality and anomaly detection with automatic quarantine or alerting before bad data feeds downstream systems?	Yes Treasure Data includes features to monitor data quality and detect anomalies in data feeds. The platform leverages AI/ML to flag unusual patterns or errors in incoming data, and it provides alerting dashboards so teams can intervene. For example, it can generate anomaly alerts or quarantine suspect data if a source suddenly deviates (preventing corrupted data from propagating). This proactive data quality management is part of Treasure Data’s “trust” capabilities, ensuring that downstream systems only receive reliable data.	https://www.treasuredata.com/press-releases/end-to-end-governance-security-and-privacy-foundation/ https://www.treasuredata.com/blog/ai-ml-data-quality/
11. Does your solution enforce role-based access controls and immutable audit logs at the field level across all regions?	Yes Treasure Data implements granular role-based access control (RBAC) and maintains detailed audit logs. Administrators can define permissions down to specific data objects or even sensitive fields (for instance, masking or restricting PII fields to authorized roles). All user and system activities are logged in immutable audit logs, providing a tamper-proof trail of who accessed or changed data. These controls apply across regions – the CDP supports multi-region data hosting while uniformly enforcing security policies and logging for compliance.	https://www.treasuredata.com/press-releases/end-to-end-governance-security-and-privacy-foundation/ https://www.treasuredata.com/blog/ai-ml-data-quality/
12. Does your solution offer built-in consent & preference management that suppresses activation instantly when permissions change (GDPR/CCPA ready)?	Yes Treasure Data’s CDP has integrated consent and preference management capabilities. It centralizes individual consent statuses and privacy preferences, and ensures that those are honored in real time for all data activations. If a customer opts out or changes consent, the platform immediately suppresses that individual’s data from any downstream marketing activation. This gives marketers a GDPR/CCPA-compliant solution out of the box – consent is stored with each unified profile and automatically enforced (no messages or exports will be sent if permission is revoked).	https://www.treasuredata.com/press-releases/end-to-end-governance-security-and-privacy-foundation/
13. Does your solution support regional data residency options (e.g., EU-only, US-only) with encryption at rest (AES-256) and in transit (TLS 1.2+)?	Yes Treasure Data allows customers to choose the region where their data is stored (for example, using AWS data centers in the EU or US to meet data residency requirements). All customer data is encrypted both at rest and in transit. The platform uses strong encryption standards – data at rest is encrypted with AES-256, and data in transit over public networks is protected via TLS (following NIST guidelines, e.g. TLS 1.2+). This ensures compliance with regional data privacy laws and high security for data storage and transfer.	https://www.treasuredata.com/technology/ https://www.treasuredata.com/security/
14. Does your solution provide reverse-ETL / audience-syndication connectors that push profiles and events to major ad, analytics, and messaging destinations natively?	Yes Yes. In addition to data ingestion, Treasure Data natively supports “activation” integrations (essentially reverse ETL). It has numerous out-of-the-box destination connectors to export audience segments or event data to all major advertising platforms, analytics tools, and messaging/CRM systems. Marketers can syndicate unified profiles and their attributes directly from Treasure Data to places like Facebook Ads, Google Ads, email service providers, analytics dashboards, etc., without needing custom development – the connectors handle the delivery.	https://www.treasuredata.com/technology/

Oct 2025 CDP Pulse Report - Treasure Data

Yes/No	Details	Source(s)
Yes	Treasure Data provides a unified profile API that external systems can query in real time. This means downstream apps can request the latest customer profile or segment info on demand via REST API calls (using profile tokens and keys). In practice, companies use this to have websites, mobile apps, or customer service tools pull the up-to-date 360° customer profile from the CDP in real time. Additionally, Treasure Data supports event streaming/webhook outputs, so systems can subscribe to real-time data updates if needed. These options ensure any downstream system can access fresh profile data on the fly.	https://docs.simondata.com/reference/simon-signal-overview
No	Treasure Data's pricing is not publicly posted in a self-serve model. The company typically provides custom quotes based on a client's data volumes and needs. There isn't a transparent published pricing page detailing costs per events or MAUs that one can readily access. Instead, pricing and any usage-based fees are discussed directly with Treasure Data sales. (Enterprise CDP pricing is generally tailored; Treasure Data does not advertise a standardized usage-based price list independent of services.)	https://www.g2.com/compare/treasure-data-customer-data-platform-vs-twilio-segment

Vendor Deep Dive



Twilio Segment

Primary Archetype: Customer Data Infrastructure / Composable

Strengths: Warehouse-native / Composable, Real-time / Edge, Governance & Residency, Pricing Transparency

Capability Lens: Data Visibility & Quality; Security, Consent & Residency; Customization & Composability

VOX Score: 74.3 / 100

Preference Coverage: 12 / 16 questions “Yes”



Talk to Twilio Segment: <https://segment.com/>



Ask VOX for Introduction

Oct 2025 CDP Pulse Report - Twilio Segment

Yes/No	Details	Source(s)
1. Does your solution provide data-observability dashboards that refresh in ≤ 60 seconds (p95) and show source-to-destination lineage for every first-party data stream?	Yes Twilio Segment offers real-time observability of the data pipeline. It provides dashboards/logs that trace each event from ingestion through any filters to its delivery in destinations, giving complete source-to-destination lineage for first-party data. This means teams can inspect what happens to every event at each stage (ingestion, processing, and outbound to each connected tool) almost immediately after events occur.	https://www.twilio.com/en-us/products/data-observability https://www.twilio.com/en-us/products/data-observability
2. Does your solution ship native inbound connectors or APIs for the top-10 commerce, marketing-automation, and service platforms without custom code or PS?	Yes Twilio Segment includes hundreds of pre-built source integrations covering major commerce, marketing automation, and service platforms. It can natively pull data from popular SaaS tools (e.g. Shopify/Stripe for e-commerce, Marketo/Mailchimp for marketing, Salesforce/Zendesk for CRM/support) via provided connectors, all without custom coding. These inbound integrations enable easy data ingestion from the top systems in those categories.	https://blog.coupler.io/adverity-alternatives-competitors/ https://blog.coupler.io/adverity-alternatives-competitors/
3. Does your solution support server-side and client-side collection for web, mobile-app, email, and in-store (POS/scanner) events in a single SDK?	Yes Twilio Segment's platform supports both client-side and server-side data collection across all major channels. It provides web libraries (JavaScript), mobile SDKs (for iOS, Android, etc.), and server APIs to capture events. This allows websites, mobile apps, backend servers, and even offline/POS systems (via API or cloud source integrations) to all send events into Segment's unified pipeline. In practice, all these events feed into one system and follow a consistent tracking plan.	https://blog.coupler.io/adverity-alternatives-competitors/ https://www.saashub.com/compare-segment-vs-webmaxy-co
4. Does your solution offer deterministic identity-resolution that stitches online cookies, MAIDs, loyalty IDs and POS receipts into one customer profile?	Yes Twilio Segment provides deterministic identity resolution via its Unify feature. It uses a unified identity graph to merge identifiers across devices and channels into a single profile. The system can ingest multiple external IDs for a user (e.g. web cookie IDs, mobile ad IDs, loyalty/member IDs, in-store transaction IDs) and will automatically match and merge events to the correct customer profile in real time. This ensures activities from online and offline touchpoints are stitched together into one unified customer record.	https://segment.com/product/unify/
5. Does your solution sync online and in-store inventory events in < 5 minutes (p95) to enable BOPIS and endless-aisle use cases?	No Twilio Segment doesn't officially publish any SLA guaranteeing sub-5-minute inventory sync times. However, according to their team, sub-5-minute sync is achievable for BOPIS and endless-aisle use cases if real-time event streaming and low-latency integrations are implemented.	https://www.saashub.com/compare-segment-vs-webmaxy-co
6. Does your solution expose decision or personalization APIs that return in < 1 second (p95) so actions fire while the session is active?	No Twilio Segment does not provide a fully native, out-of-the-box real-time decision or personalization engine with guaranteed sub-second responses. However, Segment's new Edge SDK (currently in Pilot), when integrated with edge computing platforms such as Cloudflare Workers, enables developers to build custom real-time personalization solutions that deliver content and UI adjustments within milliseconds. While this powerful capability is achievable, it still requires custom implementation rather than being a turnkey decisioning API.	https://blog.cloudflare.com/nl-nl/twilio-segment-sdk-powered-by-cloudflare-workers/
7. Does your solution include embedded predictive-merchandising models (e.g., next-best product) deployable by business users without data-science help?	Yes Twilio Segment (particularly via Twilio Engage) offers Predictive AI features that non-technical users can leverage. Marketers can enable out-of-the-box predictive models to create audiences or recommendations. For example, Segment's Predictions capability lets users predict customer behaviors (likelihood to purchase, churn, etc.) and its Recommendations feature can automatically surface products most likely to convert, all through UI controls. These features come pre-built and do not require custom data science work by the business user.	https://segment.com/solutions/ai/ https://segment.com/solutions/ai/

Oct 2025 CDP Pulse Report - Twilio Segment

Yes/No	Details	Source(s)	
8. Does your solution allow teams to upload and deploy their own ML/AI models (PMML/ONNX/Python) to score profiles or events in real time?	No	Twilio Segment doesn't support uploading custom ML models for hosted real-time scoring. It provides ways to bring your own models by training them outside Segment (for instance, in a data warehouse or ML platform) and then importing the results. Segment's approach is to ensure its data feeds your external models and to help sync model outputs back into Segment and downstream tools. However, there is no native feature to deploy a user's PMML/ONNX/Python model within the Segment infrastructure for live inference – model scoring would need to occur outside and then integrate via Segment's APIs or Reverse ETL.	https://segment.com/solutions/ai/
9. Does your solution provide an SQL workspace or direct lake-house access for segmentation, attribution, and ad-hoc analytics?	Yes	Twilio Segment enables direct access to raw customer data for analysis by integrating with data warehouses. Through its Warehouse Sync feature, Segment will load all collected events and unified profiles into the customer's own data warehouse or lake (Snowflake, BigQuery, Redshift, etc.) in near-real time. This means teams can write SQL queries on the full dataset using their warehouse or lakehouse environment for segmentation, attribution, and ad-hoc analytics. (Segment itself doesn't have a built-in SQL UI, but it makes the data available in a queryable format in your database.)	https://www.twilio.com/en-us/resource-center/data-warehouse https://www.twilio.com/en-us/resource-center/data-warehouse
10. Does your solution run continuous data-quality and anomaly detection with automatic quarantine or alerting before bad data feeds downstream systems?	Yes	Twilio Segment's Protocols functionality continuously monitors incoming data against a defined tracking plan to ensure quality. If an event does not conform (unexpected schema or values), Segment can automatically block/quarantine that "bad" data from going out to destinations. The system also supports real-time alerting – teams can set up alerts (e.g. Slack notifications) for schema violations or anomaly conditions. In practice, this means malformed or unplanned events are caught and stopped upstream, with notifications sent, so downstream tools are protected from bad data.	https://customers.twilio.com/en-us/creative-market https://www.twilio.com/en-us/products/data-observability
11. Does your solution enforce role-based access controls and immutable audit logs at the field level across all regions?	No	Twilio Segment provides role-based access control (RBAC) at the workspace and project level (with predefined roles and permissions for who can view or change settings). It also offers an Audit Trail that logs user and system activities (e.g. configuration changes) for compliance purposes (retained for 90 days). However, we found no support for fine-grained, field-level access restrictions, nor explicit mention that audit logs are immutable beyond standard retention. In essence, Segment has basic RBAC and change logging, but not at the field-by-field level across regions as described.	https://www.twilio.com/en-us/blog/insights/data/reverse-etl https://segment.com/docs/segment-app/iam/audit-trail/
12. Does your solution offer built-in consent & preference management that suppresses activation instantly when permissions change (GDPR/CCPA ready)?	Yes	Twilio Segment includes built-in privacy tools to manage user consent and data suppression in compliance with GDPR/CCPA. It allows you to enforce consent preferences across all data collection: for example, if a user opts out, Segment can immediately stop collecting or forwarding that user's data to any destinations. The Consent Management and Suppression features ensure that when a user withdraws consent or changes preferences, their data flow is halted or removed in real time across all connected systems. Segment's documentation emphasizes one-click suppression and integration with consent frameworks to honor privacy laws.	https://www.twilio.com/en-us/products/privacy-portal/gdpr https://www.twilio.com/en-us/products/privacy-portal/gdpr

Oct 2025 CDP Pulse Report - Twilio Segment

Yes/No	Details	Source(s)	
13. Does your solution support regional data residency options (e.g., EU-only, US-only) with encryption at rest (AES-256) and in transit (TLS 1.2+)?	Yes	Twilio Segment offers regional deployment options for data residency. Customers can choose to use Regional Segment in the EU or US, ensuring that data from EU users, for example, is ingested, processed, and stored on EU-based infrastructure to meet sovereignty requirements. Additionally, Twilio's security practices include strong encryption: customer data is encrypted at rest (utilizing AES-256 standards in AWS/GCP) and data in transit is protected via TLS 1.2+ for all communications. These measures align with industry best practices for security and regional compliance.	https://www.twilio.com/en-us/products/privacy-portal/gdpr https://help.twilio.com/articles/360051805394-Measures-Twilio-Takes-to-Safeguard-the-Privacy-of-Customer-Personal-Data
14. Does your solution provide reverse-ETL / audience-syndication connectors that push profiles and events to major ad, analytics, and messaging destinations natively?	Yes	Twilio Segment natively supports Reverse ETL and audience syndication. It can take unified profile data or aggregated audiences and send them out to hundreds of downstream tools. This includes major advertising platforms, analytics services, CRMs, email/SMS marketing tools, etc. For example, Segment can sync enriched customer attributes or audience memberships from your warehouse or Segment Personas into Google and Facebook Ads, marketing automation systems, or analytics dashboards. In practice, any Segment destination (450+ integrations) can receive data – meaning the platform not only collects events in, but also pushes processed profiles and events out to other applications.	https://www.twilio.com/en-us/resource-center/data-warehouse https://blog.coupler.io/adverity-alternatives-competitors/
15. Does your solution expose a unified customer-profile API/stream that any downstream system can query or subscribe to in real time?	Yes	Twilio Segment provides a unified Profile API that allows real-time retrieval of customer profiles. Downstream systems (or your custom applications) can query this API to get a customer's full profile (with identity-resolved traits and event history) in JSON, typically within a few hundred milliseconds. This supports use cases like personalizing content on the fly by pulling the latest profile data. Additionally, Segment's architecture lets you subscribe systems to data streams by adding them as destinations or using webhooks – ensuring that as events or profile updates occur, they can be delivered in real time to any tool that needs them.	https://segment.com/infrastructure/ https://segment.com/infrastructure/
16. Does your solution publish transparent, usage-based pricing (events, MTU, storage) in writing, independent of professional-services spend?	Yes	Twilio Segment's pricing model is openly documented and primarily usage-based. The plans (Free, Team, Business) are defined by monthly tracked users ("visitors") and event volumes, with clear rates for overages. For instance, the Team plan includes 10,000 MTUs and then charges a fixed rate (e.g. \$10 per additional 1,000 users) as usage scales. These prices and terms are published on Segment's website. The pricing is not tied to professional services – any consulting or support services are separate; the core pricing is based on data and user volumes and is transparently provided in writing.	https://segment.com/pricing/connections/ https://segment.com/pricing/connections/

Engagement Hubs

Vendor Deep Dive



FirstHive

Primary Archetype: Engagement Hub

Strengths: Retail Specialist; Governance

Capability Lens: In-Session Activation

VOX Score: 62.72 / 100

Preference Coverage: 10 / 16 questions “Yes”

📞 Talk to FirstHive: <https://firsthive.com>

📞 [Ask VOX for Introduction](#)

Oct 2025 CDP Pulse Report - FirstHive

Yes/No	Details	Source(s)	
1. Does your solution provide data-observability dashboards that refresh in ≤ 60 seconds (p95) and show source-to-destination lineage for every first-party data stream?	No	FirstHive's documentation does not mention dedicated data observability or lineage dashboards with a 60-second refresh. It offers real-time analytics dashboards (via "streaming APIs" for ongoing data updates) to monitor customer interactions, but nothing about source-to-destination data lineage for each stream is explicitly noted.	https://firsthive.com/analytics/
2. Does your solution ship native inbound connectors or APIs for the top-10 commerce, marketing-automation, and service platforms without custom code or PS?	Yes	FirstHive provides 300+ native connectors for popular marketing, commerce, CRM, and service tools out-of-the-box. It also offers open APIs to integrate any additional tools not already pre-built. These connectors enable data ingestion and activation without custom coding or professional services for the most common platforms.	https://firsthive.com/data-management/
3. Does your solution support server-side and client-side collection for web, mobile-app, email, and in-store (POS/scanner) events in a single SDK?	Yes	FirstHive can collect events from websites, mobile apps, emails, and in-store/POS systems. It supports client-side tagging/SDKs and server-side data feeds. For example, FirstHive's data ingestion handles internal sources like web, e-commerce, mobile apps, and retail POS, as well as SDK-based tracking and tag integrations. This allows unified event collection across online and offline channels (all feeding into the same platform).	https://firsthive.com/data-ingestion/ https://firsthive.com/data-ingestion/
4. Does your solution offer deterministic identity-resolution that stitches online cookies, MAIDs, loyalty IDs and POS receipts into one customer profile?	Yes	FirstHive's CDP performs deterministic identity resolution to unify disparate identifiers into a single customer profile. It creates a Single Customer Identity by matching data points (online cookies, device/MAIDs, loyalty program IDs, offline transactions/POS records, etc.) across sources, using both rule-based (deterministic) and machine-learning techniques. The platform continuously cross-references new data with existing profiles, enriching the unified profile in real time.	https://firsthive.com/data-management/ https://firsthive.com/single-view-of-customers/
5. Does your solution sync online and in-store inventory events in < 5 minutes (p95) to enable BOPIS and endless-aisle use cases?	No	FirstHive can ingest and unify data from both online systems and offline point-of-sale (POS) systems, which in principle enables use cases like BOPIS (Buy Online, Pick-up In Store) or "endless aisle." However, the vendor does not publicly specify a guaranteed <5-minute sync for inventory updates. There is no explicit mention of inventory event latency SLAs (though FirstHive does support real-time data processing in general).	https://firsthive.com/data-ingestion/ https://firsthive.com/data-ingestion/
6. Does your solution expose decision or personalization APIs that return in < 1 second (p95) so actions fire while the session is active?	Yes	FirstHive includes a real-time decision engine for personalization. It is designed to evaluate rules and predictive models on the fly so that personalized actions can occur during an active session. For example, the platform emphasizes "hyper-personalize every interaction...in real-time," indicating that its decisioning logic/APIs respond fast enough (sub-second) to influence live user experiences. (While exact response metrics aren't published, the intent is to deliver decisions within active session windows.)	https://firsthive.com/decision-engine/

Oct 2025 CDP Pulse Report - FirstHive

Yes/No	Details	Source(s)
Yes	FirstHive provides built-in AI/ML models for predictive analytics and merchandising. Marketers can leverage these out-of-the-box models (such as next-best product recommendations, customer propensity scores, etc.) directly through the CDP interface. These models are pre-built into the platform's "recommendation engine" and do not require data science coding – business users can deploy things like next-best-product suggestions via configuration.	https://firsthive.com/analytics/
No	There is no indication that FirstHive allows customers to bring their own custom ML models for deployment. The platform focuses on its built-in machine learning models for scoring and predictions. It does not advertise support for uploading custom PMML/ONNX/Python models into the system, so users must rely on the native AI capabilities rather than deploying proprietary algorithms.	https://firsthive.com/analytics/
Yes	FirstHive offers a data workbench and even a replica database access for advanced analysis. Users can get direct access to the unified dataset to run analyses or build models. This means analysts can use SQL or connect BI tools to FirstHive's data (the platform notes it "plays well with BI tools" for reporting and ad-hoc analytics). These capabilities enable custom segmentation, attribution analysis, and other queries on customer data outside the standard UI.	https://firsthive.com/analytics-data-organisation/ https://firsthive.com/analytics-data-organisation/
No	FirstHive's public materials do not describe an automated data quality monitoring or anomaly detection system that quarantines bad data. The platform does perform data cleansing and standardization during ingestion (to unify formats and remove inconsistencies), but it doesn't mention continuous anomaly detection or automatic blocking/alerting of bad data before it propagates. Any data quality checks appear to be manual or part of the integration process rather than a real-time autonomous feature.	https://firsthive.com/customer-data-platform-cdp/
Yes	FirstHive emphasizes strong security and governance. It implements role-based access control (RBAC) down to data field levels and keeps extensive log records of data access and changes. The platform uses enterprise-grade cloud security (e.g. multi-factor auth, audit logging, and hardware security modules for keys), helping ensure that every action is tracked. These controls are in effect regardless of region, as FirstHive's cloud deployment spans multiple geographies.	https://firsthive.com/data-management/

Oct 2025 CDP Pulse Report - FirstHive

	Yes/No	Details	Source(s)
12. Does your solution offer built-in consent & preference management that suppresses activation instantly when permissions change (GDPR/CCPA ready)?	Yes	Compliance with GDPR/CCPA is built into FirstHive. The CDP provides mechanisms for consent and preference management – for example, it creates a single source of truth for each customer's data and can honor privacy requests like data erasure (right to be forgotten). FirstHive is described as “private-by-design” and able to instantly respect updated consent settings. It has a “transparent mechanism in compliance with international data privacy laws” to suppress or allow data activation based on the customer's latest permissions.	https://firsthive.com/data-management/ https://firsthive.com/single-view-of-customers/
13. Does your solution support regional data residency options (e.g., EU-only, US-only) with encryption at rest (AES-256) and in transit (TLS 1.2+)?	Yes	FirstHive's cloud-native platform is deployed on Microsoft Azure and Amazon AWS, which means data can be hosted in various regions to meet residency requirements. While specific locales aren't enumerated on the website, customers can choose data hosting in compliance with EU-only or US-only needs. Security-wise, FirstHive leverages the cloud providers' advanced features: all data is encrypted (e.g., via AWS KMS/CloudHSM for AES-256 at rest) and transmitted securely (TLS encryption in transit). These measures ensure regional isolation and strong encryption for customer data.	https://firsthive.com/data-management/
14. Does your solution provide reverse-ETL / audience-syndication connectors that push profiles and events to major ad, analytics, and messaging destinations natively?	Yes	FirstHive not only ingests data but also has native connectors to push audiences and event data out to other platforms. It supports “reverse ETL” use cases by integrating with hundreds of marketing and adtech tools. For instance, users can directly syndicate segments to ad networks, email/SMS platforms, CRMs, or analytics tools via FirstHive's built-in connectors (300+ supported destinations). The platform even allows pushing audiences to existing external tools from its interface, enabling easy activation of customer data across channels.	https://firsthive.com/data-management/ https://firsthive.com/decision-engine/
15. Does your solution expose a unified customer-profile API/stream that any downstream system can query or subscribe to in real time?	No	FirstHive's architecture focuses on using its native integrations to share data rather than a open general-purpose profile API. The documentation does not describe a public real-time profile query API or streaming subscription for external systems. Instead, FirstHive encourages pulling or pushing data via its connectors and pre-built integrations (e.g. exporting audiences to tools, or embedding dashboards). There isn't an advertised feature where arbitrary external systems can subscribe to a live customer profile feed in real-time.	https://firsthive.com/decision-engine/
16. Does your solution publish transparent, usage-based pricing (events, MTU, storage) in writing, independent of professional-services spend?	No	FirstHive does not publish pricing on its website or collateral. There is no publicly available pricing page listing costs per event, MAU/MTU, storage, etc. Instead, pricing is handled via consultation (custom quotes). Industry reviews note that details like free trials or fixed-rate plans are not provided openly. In short, the pricing model is not transparently published in a usage-based format; you must contact FirstHive for a quote (and any services costs are likely separate).	https://www.trustradius.com/products/firsthive/pricing

Vendor Deep Dive

⊕ Cordial

Primary Archetype: Engagement Hub

Strengths: Retail Specialist

Capability Lens: In-Session Activation

VOX Score: 75.37 / 100

Preference Coverage: 12 / 16 questions “Yes”

📞 Talk to Cordial: <https://cordial.com/>

📞 [Ask VOX for Introduction](#)

Oct 2025 CDP Pulse Report - Cordial

Yes/No	Details	Source(s)	
1. Does your solution provide data-observability dashboards that refresh in ≤ 60 seconds (p95) and show source-to-destination lineage for every first-party data stream?	No	Cordial offers robust analytics dashboards for exploring customer event and messaging data, with fast performance (using Elasticsearch) cdpinstiute.org . However, it does not appear to have dedicated data observability/lineage dashboards that visualize source-to-destination data flows refreshing every 60 seconds. Cordial's focus is on real-time marketing analytics rather than specialized data pipeline lineage tracking.	https://www.cdpinstiute.org/wp-content/uploads/2022/01/AU024P-Cordial-2022-RealCDP-Package.pdf
2. Does your solution ship native inbound connectors or APIs for the top-10 commerce, marketing-automation, and service platforms without custom code or PS?	Yes	Cordial provides numerous pre-built connectors and APIs for popular commerce, marketing, and ad platforms out of the box. In fact, it comes with 20+ native integrations covering leading e-commerce, marketing automation, CRM, and digital advertising systems cdpinstiute.org . These connectors enable data ingestion and sharing with major third-party platforms without requiring custom code or professional services.	https://www.cdpinstiute.org/wp-content/uploads/2022/01/AU024P-Cordial-2022-RealCDP-Package.pdf
3. Does your solution support server-side and client-side collection for web, mobile-app, email, and in-store (POS/scanner) events in a single SDK?	Yes	Cordial's platform can collect first-party data from both client-side and server-side sources across all key channels. It offers a JavaScript listener for web, mobile SDKs for apps, and APIs to ingest server-side events. This allows capturing customer interactions on websites, mobile apps, emails, and even in-store/POS systems under one unified data platform cordial.com . All these events feed into Cordial's real-time customer data platform for use in messaging and personalization.	https://cordial.com/customer-data-platform/ https://www.cdpinstiute.org/wp-content/uploads/2022/01/AU024P-Cordial-2022-RealCDP-Package.pdf
4. Does your solution offer deterministic identity-resolution that stitches online cookies, MAIDs, loyalty IDs and POS receipts into one customer profile?	Yes	Yes. Cordial includes an identity resolution capability that unifies disparate identifiers into a single customer profile. It deterministically merges identifiers from different channels (e.g. web cookies, mobile device IDs, emails, loyalty or POS IDs) to recognize the same individual across touchpoints. Cordial's Identity+ service further enhances this by matching anonymous browser and device IDs to known users, increasing identified traffic and consolidating profiles.	https://cordial.com/identity-resolution/ https://cordial.com/identity-resolution/
5. Does your solution sync online and in-store inventory events in < 5 minutes (p95) to enable BOPIS and endless-aisle use cases?	Yes	Yes. Cordial supports real-time ingestion and processing of events (including product/inventory updates) to enable use cases like "buy online, pick up in store" and endless aisle. For example, Cordial can take an in-store inventory change or a back-in-stock update and trigger a customer message almost immediately. The platform is built for strong real-time data updates – it processes incoming events quickly and can fire off personalized messages via email, SMS, app, etc. in response in near real time.	https://cordial.com/identity-resolution/ https://www.cdpinstiute.org/wp-content/uploads/2022/01/AU024P-Cordial-2022-RealCDP-Package.pdf
6. Does your solution expose decision or personalization APIs that return in <1 second (p95) so actions fire while the session is active?	Yes	Yes. Contentstack CDP (Lytics) exposes a low-latency REST API for fetching user profiles and segment memberships in real time, with response times engineered to be under one second. This enables synchronous, in-session personalization by allowing client-side or server-side applications to retrieve user data and trigger an experience while the session is active. The system supports same-page, first-visit personalization for both anonymous and known visitors, using edge delivery for speed. In practice, this means personalization decisions are returned essentially instantly (sub-second) to be applied within the active session, yielding flicker-free, in-session experiences.	https://research.com/software/reviews/cordial https://research.com/software/reviews/cordial
7. Does your solution include embedded predictive-merchandising models (e.g., next-best product) deployable by business users without data-science help?	Yes	Yes. Cordial comes with built-in AI/ML models that marketers can use for predictive merchandising. Through Cordial Edge™ AI, the platform can automatically predict things like which product a customer is most likely to buy next or who is at risk of churn, without requiring data-science intervention. In fact, Cordial provides several out-of-the-box predictive scores (e.g. engagement score, category affinity for products) that business users can leverage to personalize campaigns.	https://cordial.com/ai/ https://www.cdpinstiute.org/wp-content/uploads/2022/01/AU024P-Cordial-2022-RealCDP-Package.pdf

Oct 2025 CDP Pulse Report - Cordial

	Yes/No	Details	Source(s)
8. Does your solution allow teams to upload and deploy their own ML/AI models (PMML/ONNX/Python) to score profiles or events in real time?	Yes	Yes. Beyond its native models, Cordial enables teams to bring their own predictive models. The platform's Architect™ data tooling lets users build custom models directly within Cordial using their data. Cordial Edge™ AI is flexible – whether using Cordial's provided models or integrating an external model, the system can incorporate those model outputs so they execute against streaming profile/event data in real time.	https://cordial.com/platform/data-platform/data-transformations/ https://cordial.com/ai/
9. Does your solution provide an SQL workspace or direct lake-house access for segmentation, attribution, and ad-hoc analytics?	Yes	Yes. Cordial avoids locking data in a silo – it provides options for direct access to your customer data for advanced analysis. Clients can establish a direct connection between Cordial and their cloud data warehouse (for example, a Snowflake data share) to query all of their first-party data in SQL. This means teams can perform custom segmentation, attribution modeling, and other ad-hoc analytics on Cordial data using SQL or BI tools, leveraging a live feed of the unified customer data.	https://cordial.com/platform-integrations/
10. Does your solution run continuous data-quality and anomaly detection with automatic quarantine or alerting before bad data feeds downstream systems?	No	Cordial includes data transformation and hygiene features (e.g. it can cleanse and standardize incoming data, including things like postal address hygiene), but it does not appear to have an automated data observability module for anomaly detection. There's no indication of continuous monitoring that would quarantine suspicious data in real-time. Data quality in Cordial is maintained through predefined transformations and validations rather than AI-driven anomaly alerts.	https://www.cdpinstitute.org/wp-content/uploads/2022/01/AU024P-Cordial-2022-RealCDP-Package.pdf
11. Does your solution enforce role-based access controls and immutable audit logs at the field level across all regions?	Yes	Yes. Cordial implements fine-grained role-based access control (RBAC) to secure data and operations. Administrators can define custom roles with specific permissions and assign them to users, ensuring each team member only accesses authorized features/data. The platform also maintains audit logs of user activities – for example, an internal User Activities log is available to track changes and actions by users. These security measures (along with SOC 2 compliance) help enforce accountability across regions and cannot be tampered with by users.	https://support.cordial.com/hc/en-us/articles/360048062092-Roles-and-Permissions https://support.cordial.com/hc/en-us/articles/360048062092-Roles-and-Permissions
12. Does your solution offer built-in consent & preference management that suppresses activation instantly when permissions change (GDPR/CCPA ready)?	Yes	Yes. Cordial provides native consent and preference management capabilities and is fully GDPR/CCPA compliant. The platform can store and honor customer consent settings and channel preferences out-of-the-box. For example, Cordial can automatically update and sync a contact's opt-in/opt-out status across all messaging – if a user withdraws consent or unsubscribes, Cordial's system will immediately stop any further activations for that user on the disallowed channels.	https://research.com/software/reviews/cordial https://cordial.com/platform/data-platform/data-transformations/
13. Does your solution support regional data residency options (e.g., EU-only, US-only) with encryption at rest (AES-256) and in transit (TLS 1.2+)?	No	Cordial adheres to enterprise-grade security (including encryption in transit via TLS and encryption at rest on AWS infrastructure) and privacy compliance, but it does not offer customer-selectable data residency regions in its standard offering. Cordial's SaaS is multi-tenant (hosted on AWS) and complies with frameworks like EU-US Privacy Shield for data transfers, rather than providing distinct EU-only or other regional datastores. (All data is protected via SOC 2 Type II controls and industry-standard encryption, but regional isolation options are not explicitly advertised.)	https://cordial.com/security/
14. Does your solution provide reverse-ETL / audience-syndication connectors that push profiles and events to major ad, analytics, and messaging destinations natively?	Yes	Yes. Cordial features native audience syndication (reverse-ETL) capabilities to push customer data to other systems. It has over 20 pre-built outbound integrations covering major advertising, analytics, and messaging platforms. For example, marketers can directly sync segments or events from Cordial to tools like Google Ads, Facebook/Meta, or other martech destinations. Additionally, Cordial's universal connector supports real-time or batch data transfer to virtually any partner system for activation.	https://www.cdpinstitute.org/wp-content/uploads/2022/01/AU024P-Cordial-2022-RealCDP-Package.pdf https://cordial.com/platform-integrations/

Oct 2025 CDP Pulse Report - Cordial

Yes/No	Details	Source(s)
Yes	<p>Yes. Cordial's unified customer profile is accessible via APIs in real time. The platform is built with an API-first approach, allowing external systems to query profile data or retrieve updates as they happen. Developers can use Cordial's RESTful API and webhook subscriptions to integrate downstream systems – for instance, a CRM or personalization engine can pull the latest profile attributes or listen for events (like segment membership changes) from Cordial in real time for immediate action.</p>	https://research.com/software/reviews/cordial
No	<p>Cordial's pricing model is not publicly transparent in a usage-based format. The company typically provides custom pricing quotes tailored to each client's volume and needs, rather than publishing a standard rate card. For example, on AWS Marketplace Cordial lists packaged pricing tiers (e.g. around \$125,000 per year for up to 125 million emails)aws.amazon.com, but these are negotiated contract plans. In general, Cordial does not openly post granular pricing for events, MAUs, or storage on its website, and pricing is handled via sales without dependence on professional services fees.</p>	https://aws.amazon.com/marketplace/pp/prodview-x2epdvgrgd65y

Vendor Deep Dive



Ometria

Primary Archetype: Engagement Hub

Strengths: Retail Specialist

Capability Lens: In-Session Activation

VOX Score: 68.41 / 100

Preference Coverage: 11 / 16 questions “Yes”

📞 Talk to Ometria: <https://ometria.com/>

📞 [Ask VOX for Introduction](#)

Oct 2025 CDP Pulse Report - Ometria

Yes/No	Details	Source(s)	
1. Does your solution provide data-observability dashboards that refresh in ≤ 60 seconds (p95) and show source-to-destination lineage for every first-party data stream?	Yes	Ometria offers a real-time dashboard for first-party data that updates continuously (within seconds) and displays live e-commerce activity. It shows incoming events (e.g. site visits, orders) along with their sources (campaign, channel, etc.), providing visibility into data flows. This real-time observability allows users to monitor data streams nearly instantaneously and see where data originates.	https://ometria.com/use-cases/data-ingestion https://support.ometria.com/hc/en-gb/articles/9518451781149-Event-stream-in-the-realtime-dashboard
2. Does your solution ship native inbound connectors or APIs for the top-10 commerce, marketing-automation, and service platforms without custom code or PS?	Yes	Ometria provides numerous out-of-the-box connectors and APIs for popular retail systems. It has native integrations for major e-commerce platforms (like Shopify, Magento, BigCommerce), marketing and loyalty tools, and more – all deployable with minimal coding. These pre-built connectors (plus REST APIs, webhooks, etc.) let businesses plug in top commerce, marketing, and service platforms without heavy custom work.	https://ometria.com/use-cases/data-ingestion https://ometria.com/use-cases/enterprise-scalability
3. Does your solution support server-side and client-side collection for web, mobile-app, email, and in-store (POS/scanner) events in a single SDK?	Yes	Ometria's platform can ingest event data from all channels both client-side and server-side. It provides a JavaScript tracker for websites, mobile SDKs (iOS/Android/React Native) for app events, and APIs for server-side or in-store/POS data. All these feed into one unified customer data platform. Web, mobile, email engagement, and offline retail events are consolidated in a single system/SDK framework.	https://ometria.com/use-cases/data-ingestion https://support.ometria.com/hc/en-gb/articles/36001858998-How-to-integrate-your-mobile-app-with-Ometria
4. Does your solution offer deterministic identity-resolution that stitches online cookies, MAIDs, loyalty IDs and POS receipts into one customer profile?	Yes	Ometria includes deterministic identity resolution that merges identifiers into unified profiles. It matches customer records across emails, mobile device IDs/cookies, loyalty or account IDs, and offline purchase records. As a result, online and offline interactions (e.g. web cookies, mobile app IDs, loyalty program info, in-store receipt data) are stitched together to build one single customer view.	https://ometria.com/use-cases/data-consolidation-and-insight https://ometria.com/use-cases/data-ingestion
5. Does your solution sync online and in-store inventory events in < 5 minutes (p95) to enable BOPIS and endless-aisle use cases?	No	Ometria is not designed as an inventory management system, and it does not guarantee sub-5-minute synchronization of inventory changes. Its data ingestion from commerce platforms typically runs on scheduled imports (e.g. an hourly sync by default for e-commerce data). Thus it isn't built to update inventory events in real-time for BOPIS or "endless aisle" scenarios without custom setup.	https://support.ometria.com/hc/en-gb/articles/5732442828317-Setting-up-an-API-based-Magento-2-integration
6. Does your solution expose decision or personalization APIs that return in < 1 second (p95) so actions fire while the session is active?	Yes	Ometria offers real-time personalization capabilities, including APIs for website personalization and decisioning. For example, its Web Personalisation API can deliver dynamic content or product recommendations on the fly. The platform is built to respond very quickly (sub-second range) so that personalized recommendations or decisions (e.g. next-best product) can be shown to users during an active session.	https://ometria.com/use-cases/experience-orchestration https://docs.ometria.com/docs/getting-started-with-website-personalisation
7. Does your solution include embedded predictive-merchandising models (e.g., next-best product) deployable by business users without data-science help?	Yes	Ometria comes with built-in AI models for merchandising and recommendations. Marketers can easily use these in campaigns via the UI. In particular, the platform provides multiple product recommendation "engines" (e.g. best next product, trending items, personalized picks) that automatically predict products a customer is likely to buy. Business users can deploy these predictive recommendations without any coding or data science required.	https://support.ometria.com/hc/en-gb/articles/360011378298-Product-recommendations-overview

Oct 2025 CDP Pulse Report - Ometria

	Yes/No	Details	Source(s)
8. Does your solution allow teams to upload and deploy their own ML/AI models (PMML/ONNX/Python) to score profiles or events in real time?	No	Ometria does not support importing or hosting custom machine learning models within the platform. Users are limited to the platform's native AI features. There is no mention in the documentation of uploading PMML/ONNX models or custom Python code for real-time scoring. Any scoring from external models would have to be done outside Ometria and then imported as data (not run internally in real time).	
9. Does your solution provide an SQL workspace or direct lake-house access for segmentation, attribution, and ad-hoc analytics?	No	Ometria does not offer a built-in SQL query interface or direct lake-house access to raw data. Instead, it enables data export and integration with external analytics environments. For advanced analysis or custom segmentation beyond the built-in reports, users can sync or export Ometria data to warehouses/BI tools (e.g. Snowflake, Tableau) and perform SQL queries there.	https://support.ometria.com/hc/en-gb/articles/17341018642717-Third-party-data-exports-and-imports
10. Does your solution run continuous data-quality and anomaly detection with automatic quarantine or alerting before bad data feeds downstream systems?	No	We found no evidence that Ometria includes an automated data quality monitoring or anomaly detection engine that quarantines bad data. The platform does provide certain alerts for specific issues (like integration sync failures or campaign errors) and logs user activities, but it lacks a general-purpose data anomaly detection feature that would automatically halt or flag suspect data before it propagates.	
11. Does your solution enforce role-based access controls and immutable audit logs at the field level across all regions?	Yes	Ometria has robust security controls with role-based access and comprehensive audit logging. User roles/permissions strictly govern access down to specific functions or data. Moreover, every significant user action is recorded in immutable audit logs – these logs cannot be altered or deleted (even admins have read-only access to them). These controls apply across regions, ensuring consistent governance and traceability.	https://ometria.com/data-processing-addendum https://support.ometria.com/hc/en-gb/articles/9883674238877-Can-an-admin-user-delete-or-edit-the-audit-logs
12. Does your solution offer built-in consent & preference management that suppresses activation instantly when permissions change (GDPR/CCPA ready)?	Yes	Ometria includes integrated consent and preference management as a core feature. It maintains a single source of truth for each contact's marketing permissions across channels. If a customer opts out or changes consent, the system immediately respects that – suppressing any further outreach to that customer until consent is restored. Ometria is designed to help clients stay compliant with GDPR/CCPA requirements.	https://ometria.com/use-cases/data-consolidation-and-insight
13. Does your solution support regional data residency options (e.g., EU-only, US-only) with encryption at rest (AES-256) and in transit (TLS 1.2+)?	Yes	Ometria can accommodate data residency needs and employs strong encryption. Clients are hosted in designated data centers (e.g. EU or US) as required. All data is protected – Ometria uses industry-standard encryption for data at rest and TLS 1.2+ for data in transit. The platform's technical safeguards (e.g. cryptographic protocols and encryption technologies) ensure compliance with regional data security standards.	https://docs.ometria.com/docs/getting-started-with-website-personalisation https://ometria.com/data-processing-addendum
14. Does your solution provide reverse-ETL / audience-syndication connectors that push profiles and events to major ad, analytics, and messaging destinations natively?	Yes	Ometria offers native connectors to sync audiences and events out to many third-party destinations. Users can easily set up connections to advertising platforms (Google Ads, Facebook), analytics and service tools (e.g. Freshdesk, Zendesk), and messaging channels (SMS via Twilio, etc.). The platform's "Connections" interface lets you push customer profiles or segments to these external systems without custom development.	https://support.ometria.com/hc/en-gb/articles/360012260218-Connections https://support.ometria.com/hc/en-gb/articles/360012260218-Connections

Oct 2025 CDP Pulse Report - Ometria

	Yes/No	Details	Source(s)
15. Does your solution expose a unified customer-profile API/stream that any downstream system can query or subscribe to in real time?	Yes	Ometria provides a unified customer profile API that allows external systems to retrieve profile data on demand. Through its RESTful Data API, clients can query individual contact profiles or lists of profiles in real time. (For example, one can GET a customer's unified profile via API, including all attributes and identifiers.) This enables downstream systems to pull the latest customer data from Ometria as needed.	https://docs.ometria.com/reference/get_contacts-collection
16. Does your solution publish transparent, usage-based pricing (events, MTU, storage) in writing, independent of professional-services spend?	No	Ometria's pricing is not publicly disclosed as a simple usage-based rate card. The company does not publish a transparent pricing page for events, monthly tracked users, storage, etc. Pricing is customized and provided via sales consultation. Prospective customers must contact Ometria for quotes; there is no openly posted usage-based pricing model separate from services.	https://research.com/software/reviews/ometria

Vendor Deep Dive



Dotdigital

Primary Archetype: Engagement Hub

Strengths: Retail Specialist

Capability Lens: In-Session Activation

VOX Score: 62.45 / 100

Preference Coverage: 10 / 16 questions “Yes”

📞 Talk to Dotdigital: <https://dotdigital.com>

📞 [Ask VOX for Introduction](#)

Oct 2025 CDP Pulse Report - Dotdigital

Yes/No	Details	Source(s)
1. Does your solution provide data-observability dashboards that refresh in ≤ 60 seconds (p95) and show source-to-destination lineage for every first-party data stream?	No Dotdigital focuses on marketing performance dashboards and real-time campaign reports, but it does not offer dedicated data observability or end-to-end lineage visualization for data pipelines. The platform's monitoring is limited to basic sync status and error logs (e.g. checking integration or data feed errors), rather than interactive lineage dashboards. Reporting in Dotdigital centers on engagement metrics and content performance, not on tracing data flow from source to destination in real time.	https://developer.dotdigital.com/docs/dotdigital-data-firehose-developer-reference https://support.dotdigital.com/en/articles/8198713-get-started-with-analytics
2. Does your solution ship native inbound connectors or APIs for the top-10 commerce, marketing-automation, and service platforms without custom code or PS?	Yes Dotdigital provides a wide range of pre-built connectors for popular e-commerce, CRM, and other systems, enabling code-free data integration. Out-of-the-box connectors exist for major commerce platforms like Adobe/Magento, Shopify, BigCommerce, Shopware, etc., as well as CRM or service platforms such as Salesforce and Microsoft Dynamics. The platform also integrates with review/loyalty tools (e.g. Trustpilot, Yotpo) and offers an open API and iPaaS connectors (e.g. Zapier, Power Automate) to plug into other top marketing and service apps. These native integrations cover the top marketing automation and commerce systems, avoiding the need for custom development or professional services.	https://support.dotdigital.com/en/articles/9729831-integrations-by-category https://support.dotdigital.com/en/articles/9729831-integrations-by-category
3. Does your solution support server-side and client-side collection for web, mobile-app, email, and in-store (POS/scanner) events in a single SDK?	No Dotdigital supports both client-side and server-side data collection, but not through one unified SDK for all channels. For web, it provides the Dotdigital Tag (JavaScript library) to track page views, logins, cart updates, etc., and identify users by email/phone on the client side. For mobile apps, it offers native iOS/Android SDKs focused on push notifications (requiring an email/ID to tie the device to a profile). Email engagement is tracked server-side within the platform automatically. In-store/POS events can be ingested via APIs or file import into contact records (e.g. as Insight data), and through connectors (for example, the NetSuite POS integration syncs offline purchase data). However, there isn't a single combined SDK that captures web, mobile, email, and point-of-sale events together; instead, Dotdigital uses separate tools (APIs, JS tag, SDKs) for different channels.	https://developer.dotdigital.com/docs/dotdigital-tag-api-methods https://developer.dotdigital.com/docs/ios-sdk-new
4. Does your solution offer deterministic identity-resolution that stitches online cookies, MAIDs, loyalty IDs and POS receipts into one customer profile?	Yes Dotdigital builds a Single Customer View (SCV) by deterministically unifying identifiers from multiple touchpoints into one profile. The platform's tracking tag can identify a visitor (e.g. when they log in or subscribe) by email or phone number, linking browser cookies to that known contact record. Similarly, mobile app users are tied to their profile via captured email or mobile ID. Loyalty or offline purchase data (like POS transactions) can be imported or synced into the same contact record using matching identifiers (email, customer ID, etc.), so all online and offline interactions are stitched together under one customer profile. In practice, Dotdigital uses explicit keys (rather than probabilistic matching) – when the same identifier (such as email or loyalty ID) is present, events and attributes are merged into a single unified profile. This deterministic approach allows the platform to recognize the same individual across cookies, devices, and offline records, consolidating data into a single view for personalization.	https://developer.dotdigital.com/docs/dotdigital-tag-api-methods https://dotdigital.com/blog/what-is-a-cxdp-customer-experience-and-data/
5. Does your solution sync online and in-store inventory events in < 5 minutes (p95) to enable BOPIS and endless-aisle use cases?	Yes Dotdigital can capture inventory updates from online and in-store systems within minutes via its integrations. The platform's commerce connectors (for platforms like Shopify, Shopline, WooCommerce, etc.) typically run data synchronizations on a ~5-minute cycle, ensuring that changes in product stock or orders are quickly reflected. For example, Dotdigital's integration with NetSuite (which covers POS) and other retail systems will import product catalog stock levels and order events into the platform almost in real time. This near-real-time sync capability supports use cases like Buy-Online-Pickup-In-Store and "endless aisle" availability, by keeping inventory data up-to-date across channels. (In practice, sync frequency may vary with data volume, but ~5-minute updates are the typical service level.)	https://support.dotdigital.com/en/articles/8948634-understanding-the-sync-process-from-shopline-to-dotdigital https://dotdigital.com/blog/what-is-new-hyper-personalization-and-optimized-data/
6. Does your solution expose decision or personalization APIs that return in < 1 second (p95) so actions fire while the session is active?	No Dotdigital does not provide a stand-alone real-time decision engine API for external queries; its personalization and decisioning mainly occur within its own automation and messaging workflows. The platform focuses on trigger-based communications (emails, SMS, etc.) and on-page personalization via integrated partners rather than an open decision API. For on-site web personalization, Dotdigital leverages partnerships like Fresh Relevance – e.g. marketers can drag-and-drop Fresh Relevance Smart Blocks to deliver real-time personalized content (product recommendations, countdown timers, etc.) on websites or in emails. While Dotdigital's APIs and webhooks enable real-time messaging or data sync (ensuring minimal delay in triggering actions during an active session), the solution doesn't offer a native "decisioning" API that, for instance, returns next-best-action or content within <1s for arbitrary external personalization queries. Instead, real-time personalization is achieved through its built-in automation triggers and integrations, not a standalone sub-second response API.	https://dotdigital.com/blog/what-is-new-hyper-personalization-and-optimized-data/ https://dotdigital.com/api/

Oct 2025 CDP Pulse Report - Dotdigital

	Yes/No	Details	Source(s)
7. Does your solution include embedded predictive-merchandising models (e.g., next-best product) deployable by business users without data-science help?	Yes	<p>Dotdigital provides built-in predictive analytics for merchandising, allowing marketers to easily use these models in campaigns. The platform's Product Recommendations feature includes multiple AI-driven recommendation types – for example, "Best next" product recommendations (to predict the next likely purchase) and "Lookalikes" (collaborative filtering that finds products a customer might want based on similar shoppers) – which are powered by machine learning but exposed in a user-friendly way. Marketers can configure recommendation blocks through a visual builder (choosing the model type and filters) and then drag-and-drop these blocks into emails or pages via Dotdigital's editor. These models (like also-bought, trending products, etc.) are pre-trained/automated; no data science or coding is required to deploy them. Business users simply select the desired predictive model in the interface, and the system will automatically generate personalized product suggestions for each customer. (The ML models refresh their recommendations periodically – e.g. daily or weekly – in the background.)</p>	https://support.dotdigital.com/en/articles/8199208-product-recommendations-overview https://support.dotdigital.com/en/articles/8199208-product-recommendations-overview
8. Does your solution allow teams to upload and deploy their own ML/AI models (PMML/ONNX/Python) to score profiles or events in real time?	No	<p>Dotdigital does not support importing custom machine learning models for real-time scoring. The platform's AI capabilities are limited to built-in features (such as its WinstonAI engine and the out-of-the-box predictive models for product recommendations and segmentation). There is no facility for a user to upload a custom PMML/ONNX model or run custom Python/R scripts within Dotdigital. Instead, the solution focuses on providing its own AI-driven functionalities (e.g. pre-built intelligent segments, send time optimization, product recommendations, etc.) for users to leverage. Any custom advanced modeling would have to be done outside the platform and then imported as data (for example, scoring contacts via an external data science workflow and syncing the results via the API). In summary, Dotdigital's AI is "closed-box" – offering built-in ML features but no provision for deploying your own algorithms in the system.</p>	https://dotdigital.com/artificial-intelligence/ https://dotdigital.com/artificial-intelligence/
9. Does your solution provide an SQL workspace or direct lake-house access for segmentation, attribution, and ad-hoc analytics?	No	<p>Dotdigital does not offer a native SQL query console or direct access to its underlying data warehouse for arbitrary analytics. The platform instead provides a visual segmentation UI and built-in reporting dashboards for standard needs. For more advanced or ad-hoc analysis, Dotdigital enables data export to external analytics environments: for instance, it offers a Data Firehose integration that can continuously sync your contact and event data to a Snowflake database or other data store. This allows teams to analyze data using their own BI or data lake, but it happens outside of Dotdigital. Within the platform, users cannot run custom SQL queries directly against customer data. Any custom segmentation or attribution analysis beyond the provided tools would be done by extracting data (via exports or the open API) and querying it in an external environment. Dotdigital's approach is to integrate with your analytics stack rather than act as a full SQL analytics workspace itself.</p>	https://support.dotdigital.com/en/articles/9729831-integrations-by-category https://support.dotdigital.com/en/articles/9729831-integrations-by-category
10. Does your solution run continuous data-quality and anomaly detection with automatic quarantine or alerting before bad data feeds downstream systems?	Yes	<p>Dotdigital implements some automated data-quality protections. Notably, it has a "Watchdog" service that continuously monitors contact data imports for anomalies – if a contact upload deviates significantly from usual patterns (for example, a sudden surge of suspicious addresses), the system will automatically halt the import to prevent potentially bad data (such as spam traps) from entering the platform. This proactive quarantine mechanism stops abnormal data at the point of entry. Additionally, Dotdigital provides error logging and notifications for integration syncs (administrators are alerted if a data sync fails or encounters errors), ensuring bad or incomplete data transfers are flagged early. While the platform may not have extensive AI-driven anomaly detection on every data stream, it does have built-in safeguards and alerts to maintain data quality – including automatically preventing questionable data from flowing downstream until reviewed.</p>	https://dotdigital.com/trust-center/technical-and-organizational-security-measures/ https://developer.dotdigital.com/docs/dotdigital-data-firehose-developer-reference
11. Does your solution enforce role-based access controls and immutable audit logs at the field level across all regions?	Yes	<p>Dotdigital supports granular role-based access control (RBAC) and keeps audit logs of account activities. Administrators can define user roles/permissions to restrict access to specific functions or data (for example, limit who can export contacts, who can manage accounts, etc.). Every user in the platform is assigned permissions according to their role, and these controls apply across all regions (Dotdigital instances in the EU, US, etc. all offer the same permission framework). In terms of auditing, Dotdigital maintains event logs for account activities – the system has event logging and account auditing in place. This means changes or key actions (imports, exports, user actions) are recorded for compliance and security reviews. The audit logs are not end-user editable (immutable), ensuring a reliable trail. While the RBAC is robust (though not literally at the single-field data-field level, but rather at feature/module level), and the audit logging covers user and system events, these features together provide enterprise-grade governance and accountability across regions.</p>	https://dotdigital.com/trust-center/technical-and-organizational-security-measures/ https://dotdigital.com/trust-center/technical-and-organizational-security-measures/

Oct 2025 CDP Pulse Report -Dotdigital

	Yes/No	Details	Source(s)
12. Does your solution offer built-in consent & preference management that suppresses activation instantly when permissions change (GDPR/CCPA ready)?	Yes	<p>Dotdigital includes consent and preference management features out of the box. It provides tools like subscription forms and preference centers that allow contacts to specify their marketing preferences (topics, channels, frequency) and to grant or withdraw consent. These preferences are stored on the contact record and are honored by the platform in real time – for example, if a user opts out or changes consent, that contact will be immediately suppressed from outbound campaigns as required. Dotdigital is designed with GDPR/CCPA compliance in mind: it tracks consent statuses and offers mechanisms like instant unsubscribe/opt-out, Do Not Contact flags, and audit trails of consent changes. The Marketing Preferences module lets marketers collect and respect zero-party data about customer preferences easily, and any change (such as a customer revoking email consent) will automatically stop messages to that individual. In short, consent changes propagate instantly throughout the platform to ensure no unauthorized communications are sent, supporting compliance with privacy regulations.</p>	https://dotdigital.com/customer-data/ https://dotdigital.com/customer-data/
13. Does your solution support regional data residency options (e.g., EU-only, US-only) with encryption at rest (AES-256) and in transit (TLS 1.2+)?	Yes	<p>Dotdigital offers secure data handling and regional hosting to meet data residency needs. The platform uses TLS 1.2+ encryption for data in transit and AES-256 encryption for data at rest to protect customer data. Moreover, Dotdigital has multiple regional infrastructures (for example, EU and US data centers); clients can have their data stored in a region that aligns with their residency or compliance requirements. The company's policies state that backup data remains within the same geographic region as the account, ensuring EU data stays in EU, US data stays in US, etc., unless otherwise arranged. This regional containment, combined with strong encryption and ISO 27001-based security controls, means the solution can satisfy EU-only or US-only data residency demands. (Customers typically choose the regional instance during onboarding, and Dotdigital ensures all data processing and storage for that account occurs in that locale.)</p>	https://dotdigital.com/trust-center/technical-and-organizational-security-measures/ https://dotdigital.com/trust-center/technical-and-organizational-security-measures/
14. Does your solution provide reverse-ETL / audience-syndication connectors that push profiles and events to major ad, analytics, and messaging destinations natively?	Yes	<p>Dotdigital can act as a hub to syndicate audience data out to other marketing and analytics platforms. It includes Ad Audience connectors that natively push segments to major ad networks – for instance, you can automatically sync contact lists (audiences) from Dotdigital to Facebook Custom Audiences, Google Ads, and even TikTok Ads for retargeting campaigns. This is achieved through program nodes and native integrations that keep those external audiences updated. In addition, Dotdigital supports sending event or segment data to other systems via its open API and webhooks, and it integrates with customer data infrastructure (like Snowflake or analytics tools via the Data Firehose) to feed data into analytics platforms. For messaging, Dotdigital itself is an execution platform (email, SMS, push, etc.), but it also provides connectors (or uses iPaaS solutions) to send data to downstream channels like CRM or service platforms. The key ad and social channels are covered natively – e.g. automated list sync for Facebook/Instagram and Google – which qualifies as built-in reverse ETL for audience activation.</p>	https://dotdigital.com/retargeting/ https://dotdigital.com/retargeting/
15. Does your solution expose a unified customer-profile API/stream that any downstream system can query or subscribe to in real time?	Yes	<p>Dotdigital exposes RESTful APIs to access customer profiles and data in real time, as well as mechanisms to subscribe to events. The platform's API includes endpoints to retrieve contacts (by ID or email) along with all their data fields, preferences, and Insight data, effectively serving as a customer-profile API. This allows downstream systems to query the Dotdigital contact database on demand. Additionally, Dotdigital supports webhooks and its "Data Firehose" for streaming data out – for example, you can set up webhooks to push events (like contact added, email bounced, etc.) to external endpoints in real time, or use Data Firehose to continuously sync contacts and their engagement events to an external data store. Dotdigital's open integration approach (including a rich developer hub) ensures that any external system can programmatically get a unified view of a customer or subscribe to updates. In practice, this means you can treat Dotdigital as a source of truth for profiles – querying it for profile data or receiving near-real-time streams of events to feed other tools in your stack.</p>	https://dotdigital.com/api/ https://developer.dotdigital.com/docs/dotdigital-data-firehose-developer-reference
16. Does your solution publish transparent, usage-based pricing (events, MTU, storage) in writing, independent of professional-services spend?	No	<p>Dotdigital does not publicly disclose usage-based pricing on its website; instead, it follows a custom quotation model. Pricing is typically based on factors like the size of your contact database, messaging volume, and feature packages, but these rates are provided via sales consultation rather than a transparent pricing page. The official Dotdigital pricing page invites users to "request a customized pricing plan" and indicates that plans are tailored to the brand's needs (contacts, send volume, pageviews, etc.), rather than showing fixed per-unit costs. There is no published rate card for events, MAUs/MTUs, or data storage that one can self-serve; all pricing details are given in writing through proposals from Dotdigital's sales team. This means pricing is not fully transparent or usage-metred in a self-service way – it's negotiated, and generally separate from any professional services (which are optional), but the exact costs are not publicly listed.</p>	https://dotdigital.com/pricing/ https://dotdigital.com/pricing/

Vendor Deep Dive



Listrak

Primary Archetype: Engagement Hub

Strengths: Real-time/Edge

Capability Lens: In-Session Activation

VOX Score: 62.32 / 100

Preference Coverage: 10 / 16 questions “Yes”

📞 Talk to Listrak: <https://www.listrak.com/>

📞 [Ask VOX for Introduction](#)

Oct 2025 CDP Pulse Report - Listrak

	Yes/No	Details	Source(s)
1. Does your solution provide data-observability dashboards that refresh in ≤ 60 seconds (p95) and show source-to-destination lineage for every first-party data stream?	No	Listrak's platform includes real-time marketing insights and dashboards (e.g. "70+ actionable data visualizations") for campaign performance and customer behavior, but it does not mention dedicated data observability or lineage tracking for data pipelines. There is no evidence of a 60-second refresh, source-to-destination lineage dashboard in Listrak's documentation – its focus is on marketing metrics and customer insights rather than low-level data stream monitoring.	https://www.listrak.com/blog/cdp-or-crm-what-do-retailers-really-need
2. Does your solution ship native inbound connectors or APIs for the top-10 commerce, marketing-automation, and service platforms without custom code or PS?	Yes	Yes. Listrak offers many pre-built integrations to ingest data from leading e-commerce, CRM, loyalty, and other systems with minimal effort. For example, it provides native connectors for Shopify Plus, Adobe/Magento, BigCommerce, Salesforce Data 360, etc. Listrak's integration ecosystem covers top commerce platforms and tech partners, enabling real-time data flow and identity resolution without requiring custom coding or professional services.	https://www.listrak.com/platform/integrations https://www.listrak.com/platform/integrations
3. Does your solution support server-side and client-side collection for web, mobile-app, email, and in-store (POS/scanner) events in a single SDK?	Yes	Yes. Listrak supports both client-side and server-side data collection across channels via one unified platform. It provides a JavaScript integration for website tracking/personalization and RESTful APIs for server-side data ingestion (covering contacts, custom events, order and product data, etc.). Mobile app engagement (push notifications) can be integrated via its Mobile API, and offline/in-store events (like POS transactions) can be imported via its data APIs – all feeding into the same customer profile.	https://www.listrak.com/platform/integrations https://www.listrak.com/learn/developers
4. Does your solution offer deterministic identity-resolution that stitches online cookies, MAIDs, loyalty IDs and POS receipts into one customer profile?	Yes	Yes. Listrak provides industry-leading identity resolution that unifies customer identities across devices, channels, and data sources into a single profile. It continuously stitches together online behaviors (web cookies, device IDs) with offline and transactional identifiers (like loyalty program IDs or in-store purchase data), giving an accurate omnichannel view of each individual. This deterministic identity resolution engine connects online and offline behavior to one person for 1:1 personalization.	https://www.listrak.com/platform/data-platform https://www.listrak.com/platform/data-platform
5. Does your solution sync online and in-store inventory events in < 5 minutes (p95) to enable BOPIS and endless-aisle use cases?	Yes	Yes. Listrak's platform integrates product and inventory data in real time from e-commerce and in-store systems. It captures inventory changes and product feeds continuously and can trigger responsive campaigns (e.g. back-in-stock alerts, low-stock notifications) within minutes. This real-time data synchronization enables retail use cases like Buy Online Pickup In Store (BOPIS) and endless-aisle, as inventory updates are available almost immediately for personalization and messaging workflows.	https://www.listrak.com/platform/integrations https://www.listrak.com/platform/integrations
6. Does your solution expose decision or personalization APIs that return in < 1 second (p95) so actions fire while the session is active?	Yes	Yes. Listrak supports real-time, in-session personalization. Its platform includes on-site personalization and decisioning capabilities (e.g. dynamic popups, product recommendations) that operate with immediate context. Marketers can set up real-time targeting rules and the system evaluates conditions instantly (page scrolls, clicks, etc.) to personalize content for the user on the fly. These personalization features are designed to respond within active sessions, effectively delivering sub-second decisioning on the website.	https://www.listrak.com/platform/experience-builder#Onsite-Popups-Personalization https://www.listrak.com/platform/experience-builder#Onsite-Popups-Personalization
7. Does your solution include embedded predictive-merchandising models (e.g., next-best product) deployable by business users without data-science help?	Yes	Yes. Listrak comes with built-in AI and predictive models for merchandising and customer targeting that non-technical users can leverage. It offers an AI-driven Product Recommendations engine (to suggest next-best products like top sellers, affinity-based items, etc.) which marketers can configure visually. It also provides predictive analytics such as likelihood-to-purchase or churn scores for segmentation. All these models are embedded in the platform and usable directly by marketers without needing a data scientist.	https://www.listrak.com/platform/listrak-intelligence https://www.listrak.com/platform/listrak-intelligence

Oct 2025 CDP Pulse Report - Listrak

Yes/No	Details	Source(s)	
8. Does your solution allow teams to upload and deploy their own ML/AI models (PMML/ONNX/Python) to score profiles or events in real time?	No	Listrak does not support importing custom ML models by clients. The platform's machine learning capabilities are provided out-of-the-box (e.g. its own predictive algorithms and AI for recommendations and scoring). There is no mention of uploading or deploying external models (PMML, ONNX, Python) into Listrak's environment. Users are expected to use Listrak's built-in AI intelligence features, rather than bringing their own machine learning models for real-time scoring.	https://www.listrak.com/blog/cdp-or-crm-what-do-retailers-really-need
9. Does your solution provide an SQL workspace or direct lake-house access for segmentation, attribution, and ad-hoc analytics?	No	No. Listrak does not offer a user SQL workspace or direct access to a data lake. Instead, it provides advanced segmentation tools and pre-built analytics dashboards through its interface. Marketers can segment audiences and view performance metrics using Listrak's UI (with features like unified customer profiles and dynamic segments), but ad-hoc analysis is done via the platform's reporting suite and visualization tools, not via direct SQL queries on the raw data or external lake-house access.	https://www.listrak.com/blog/cdp-or-crm-what-do-retailers-really-need https://www.listrak.com/platform/reporting-analytics
10. Does your solution run continuous data-quality and anomaly detection with automatic quarantine or alerting before bad data feeds downstream systems?	No	No. There is no indication that Listrak has an automated data-quality monitoring or anomaly detection system that quarantines bad data. Ensuring data quality in Listrak tends to be a matter of proper integration and validation by users/administrators. The platform documentation suggests marketers are responsible for identifying data anomalies or inconsistencies as they analyze campaign results. In practice, Listrak does not advertise any built-in continuous anomaly alerts or quarantines on incoming data feeds.	https://www.listrak.com/blog/how-data-management-is-moving-from-it-to-marketing-teams https://www.listrak.com/blog/how-data-management-is-moving-from-it-to-marketing-teams
11. Does your solution enforce role-based access controls and immutable audit logs at the field level across all regions?	Yes	Yes. Listrak provides enterprise-grade security controls including role-based access control (RBAC) and activity auditing. Administrators can assign built-in user roles and permissions to limit access on a least-privilege principle. The platform also maintains audit logs of user activities in the system to support accountability. These security features apply across Listrak's environments (with SSO, 2FA, etc.), ensuring that all user actions are tracked and access to sensitive data is appropriately restricted by role.	https://www.listrak.com/privacy-and-terms/security-policy https://www.listrak.com/privacy-and-terms/security-policy
12. Does your solution offer built-in consent & preference management that suppresses activation instantly when permissions change (GDPR/CCPA ready)?	Yes	Yes. Listrak includes built-in consent and preference management at the contact level. The platform tracks subscription status and consent for each individual across channels (email, SMS, etc.). If a customer opts-out or changes preferences, Listrak immediately updates their profile and honors those changes in all activations (ensuring no messages are sent if consent is withdrawn). Listrak markets itself as GDPR/CCPA-ready, with mechanisms to manage opt-ins/opt-outs and compliance flags in real time.	https://www.listrak.com/platform/data-platform
13. Does your solution support regional data residency options (e.g., EU-only, US-only) with encryption at rest (AES-256) and in transit (TLS 1.2+)?	No	Listrak meets security standards by ensuring encryption of data at rest and in transit (using strong encryption and TLS 1.2+ for all communications). However, it does not explicitly offer configurable regional data residency on a per-customer basis. Listrak's infrastructure is compliant with privacy regulations (SOC 2, GDPR, etc.), but customer data is generally hosted in its secure data centers without an option for EU-only or other region-specific storage partitions mentioned in public documentation.	https://www.listrak.com/marketer-challenges/i-need-to-find-a-marketing-platform-that-will-scale-with-my-growing-business https://www.listrak.com/marketer-challenges/i-need-to-find-a-marketing-platform-that-will-scale-with-my-growing-business
14. Does your solution provide reverse-ETL / audience-syndication connectors that push profiles and events to major ad, analytics, and messaging destinations natively?	Yes	Yes. Listrak offers audience syndication (Listrak Exchange) to push customer data to external marketing and ad platforms. It has native connectors to sync segments and events directly into major ad networks like Meta/Facebook and Google Ads (and others). This allows marketers to automatically update custom audiences for advertising without manual export. Listrak's platform also integrates with various analytics and marketing tools, enabling profile data and event triggers to be shared across systems in real time.	https://www.listrak.com/platform/paid-media-social-audiences https://www.listrak.com/platform/paid-media-social-audiences

Oct 2025 CDP Pulse Report - Listrak

Yes/No	Details	Source(s)
15. Does your solution expose a unified customer-profile API/stream that any downstream system can query or subscribe to in real time?	Yes. Listrak provides RESTful APIs that give access to its unified customer profiles and events for external systems. Developers can use Listrak's APIs to retrieve or update contact data, query engagement metrics, and import/export customer information programmatically. While Listrak doesn't advertise a continuous "stream" subscription, it supports real-time data exchange – for example, external systems can query profile data on-demand via API and use webhooks or automated syncs to get updates as they occur.	https://www.listrak.com/platform/integrations https://www.listrak.com/learn/developers
16. Does your solution publish transparent, usage-based pricing (events, MTU, storage) in writing, independent of professional-services spend?	Listrak's pricing is not publicly transparent or usage-based on its website. The company typically provides customized quotes based on the client's contact volume and features needed, rather than publishing a standard rate card. External comparisons note that Listrak requires contacting their sales team for pricing details. (For reference, one source cites a starting price around \$150/month for 5,000 contacts, but actual costs vary.) There is no openly posted usage-tier pricing independent of negotiations or services.	https://www.groovecommerce.com/ecommerce-resources/listrak-vs-klaviyo/ https://www.groovecommerce.com/ecommerce-resources/listrak-vs-klaviyo/

Vendor Deep Dive



Oracle Unity CDP

Primary Archetype: Engagement Hub

Strengths: Real-time/Edge

Capability Lens: In-Session Activation

VOX Score: 87.75 / 100

Preference Coverage: 14 / 16 questions “Yes”

📞 Talk to Oracle Unity CDP: <https://www.oracle.com>

📞 Ask VOX for Introduction

Oct 2025 CDP Pulse Report - Oracle Unity CDP

Yes/No	Details	Source(s)	
1. Does your solution provide data-observability dashboards that refresh in ≤ 60 seconds (p95) and show source-to-destination lineage for every first-party data stream?	No	Oracle Unity CDP focuses on consolidating and cleansing customer data in one place, but it does not advertise real-time data observability dashboards with sub-minute refresh or full source-to-target lineage visualization. While Unity ensures enterprise data quality by cleansing and preparing ingested data, there is no indication of a built-in lineage dashboard that updates every 60 seconds.	https://www.oracle.com/cx/customer-data-platform/
2. Does your solution ship native inbound connectors or APIs for the top-10 commerce, marketing-automation, and service platforms without custom code or PS?	Yes	Oracle Unity offers many pre-built connectors across major commerce, marketing, and service systems. It natively integrates with leading commerce platforms (e.g. Oracle Commerce, Salesforce Data 360, SAP Hybris, Shopify), marketing automation and email tools (Oracle Responsys, Oracle Eloqua, Adobe Campaign, Marketo, Mailchimp, Salesforce Marketing Cloud, Twilio SendGrid), and customer service apps (Oracle Service, Salesforce Service Cloud, ServiceNow), among others, all without custom development.	https://www.oracle.com/cx/customer-data-platform/unity-integrations/
3. Does your solution support server-side and client-side collection for web, mobile-app, email, and in-store (POS/scanner) events in a single SDK?	Yes	Oracle's CDP supports both client-side and server-side data collection across channels. Oracle provides a web tagging library (Oracle CX Tag) and a Mobile SDK for apps to capture behavioral events in real time. It also offers server-side ingestion APIs for sending events (e.g. POS transactions, profile updates) directly from source systems. Unity ingests web, mobile, email, and offline events (such as in-store/POS data) in near real time, unifying them into the customer profile.	https://docs.oracle.com/en/cloud/saas/cx-unity/cx-unity-user/Help/GetStarted/Overview_CXUnity.htm https://docs.oracle.com/en/cloud/saas/cx-unity/cx-unity-user/Help/GetStarted/Overview_CXUnity.htm
4. Does your solution offer deterministic identity-resolution that stitches online cookies, MAIDs, loyalty IDs and POS receipts into one customer profile?	Yes	Oracle Unity performs deterministic identity resolution to unify customer records. It uses configurable rules to deduplicate and merge identifiers from multiple sources (devices, channels, loyalty systems, in-store transactions, etc.) into a single customer profile. The platform's "Flexible ID resolution" and scheduled ID graph jobs resolve identities across cookies/devices, loyalty or CRM IDs, and other keys to build one unified 360° customer record.	https://www.oracle.com/cx/customer-data-platform/ https://docs.oracle.com/en/cloud/saas/cx-unity/cx-unity-user/Help/GetStarted/Overview_CXUnity.htm
5. Does your solution sync online and in-store inventory events in < 5 minutes (p95) to enable BOPIS and endless-aisle use cases?	Yes	Oracle Unity can handle near real-time data updates, supporting scenarios like BOPIS. The platform provides a streaming ingestion API with an expected latency under one minute from event capture to Unity's database. This means time-sensitive data (e.g. inventory changes or store stock events) can be ingested and made available within minutes. While Oracle's documentation doesn't explicitly call out "BOPIS" by name, Unity's ability to ingest and update profiles or datasets in well under 5 minutes enables those omnichannel inventory use cases.	https://docs.oracle.com/en/cloud/saas/cx-unity/cx-unity-develop/docs/api/stream/real-time-api.htm
6. Does your solution expose decision or personalization APIs that return in < 1 second (p95) so actions fire while the session is active?	Yes	Oracle's platform supports real-time personalization and decisioning. Unity's architecture (with embedded Oracle database AI and integration with Oracle Infinity and Maxymiser) enables instantaneous next-best-action recommendations and trigger responses during live sessions. For example, Oracle Unity can gather customer signals in real time and deliver the "right message at the right time" across channels; its native AI models provide real-time recommendations to personalize the experience in the moment, with sub-second response times for in-session actions.	https://www.oracle.com/cx/customer-data-platform/ https://www.oracle.com/cx/customer-data-platform/
7. Does your solution include embedded predictive-merchandising models (e.g., next-best product) deployable by business users without data-science help?	Yes	Oracle Unity comes with extensive out-of-the-box AI modeling. It provides more than 27 prebuilt AI/ML models for customer analytics – including propensity scores, churn likelihood, product recommendations and next-best offers – which can be utilized by marketers directly. These "Intelligent Attributes" and predictive models are built into Unity's UI (the Intelligence Workbench), allowing business users to apply next-best-product recommendations and other predictive merchandising tactics without needing a data scientist.	https://www.oracle.com/cx/customer-data-platform/

Oct 2025 CDP Pulse Report - Oracle Unity CDP

	Yes/No	Details	Source(s)
8. Does your solution allow teams to upload and deploy their own ML/AI models (PMML/ONNX/Python) to score profiles or events in real time?	Yes	Oracle Unity supports a “bring your own model” capability. Users can import custom machine learning models into Unity’s environment to generate scores or predictions on their customer data. This allows organizations to deploy proprietary models (built in Python or exported to formats like PMML) within Unity and have the platform run those models natively (inside the Oracle infrastructure) to score profiles or events in real time, without external data movement.	https://www.oracle.com/cx/customer-data-platform/
9. Does your solution provide an SQL workspace or direct lake-house access for segmentation, attribution, and ad-hoc analytics?	Yes	Oracle Unity offers robust analytics access on unified data. The platform integrates natively with Oracle Analytics Cloud, enabling users to perform custom analyses and visualizations directly on the CDP data. Unity’s underlying data is stored in an Oracle Autonomous Data Warehouse, so organizations can query the customer data via BI tools or SQL if needed. This means marketers and analysts can run ad-hoc segmentation, attribution queries, or join to external data through Oracle’s analytics integration or by exporting data to their own data lake/warehouse environment.	https://www.oracle.com/cx/customer-data-platform/ https://www.oracle.com/news/announcement/ocw24-oracle-unity-cdp-helps-organizations-leverage-ai-to-grow-revenue-2024-09-11/
10. Does your solution run continuous data-quality and anomaly detection with automatic quarantine or alerting before bad data feeds downstream systems?	No	Oracle Unity ensures high data quality by cleansing and standardizing incoming data, but it does not advertise automated anomaly detection or quarantining of bad data in real-time. There are data management processes (like validation rules during ingest and scheduled data quality jobs), but no built-in AI-driven anomaly monitoring feature is publicly described. Data issues would typically be addressed through Oracle’s data integration tooling or custom configurations rather than an automatic quarantine within Unity.	https://www.oracle.com/cx/customer-data-platform/
11. Does your solution enforce role-based access controls and immutable audit logs at the field level across all regions?	Yes	Oracle Unity includes robust data governance and security controls. It supports role-based access control, allowing administrators to restrict data access at granular levels (by data object or field) and manage permissions across different business units or regions. Oracle also meets enterprise security standards (SOC 2, ISO 27001, HIPAA attestations), which implies the platform maintains detailed audit logs of data access and changes. Field-level auditing specifics aren’t public, but compliance certifications indicate that Unity provides immutable logging and traceability required for regulatory compliance.	https://www.oracle.com/cx/customer-data-platform/
12. Does your solution offer built-in consent & preference management that suppresses activation instantly when permissions change (GDPR/CCPA ready)?	Yes	Yes – Oracle Unity has native support for consent and preference management. It integrates with leading Consent Management Platforms (like OneTrust) so that any update to a customer’s consent or opt-in/opt-out status is immediately reflected in the Unity profile. Consent attributes are part of the Unity data model and are honored in all segmentations and activations, ensuring that if a user withdraws consent, that customer’s data will instantly be excluded from downstream marketing actions. Oracle Unity is designed to help clients comply with GDPR, CCPA and similar privacy regulations.	https://www.oracle.com/cx/customer-data-platform/
13. Does your solution support regional data residency options (e.g., EU-only, US-only) with encryption at rest (AES-256) and in transit (TLS 1.2+)?	Yes	Oracle’s CDP is deployed on Oracle Cloud Infrastructure, which provides multiple regional data center options (EU, US, etc.) to meet data residency requirements. Customers can choose to keep their Unity data in specific regions to comply with local laws. All data in Oracle Unity is protected by enterprise-grade security: data is encrypted at rest (Oracle uses AES-256 encryption for stored data) and in transit (TLS 1.2+ for data transfer). These security measures and regional deployment options ensure compliance with strict data residency and security standards.	https://www.oracle.com/a/ocom/docs/applications/cx/oracle-unit-cdp-datasheet.pdf
14. Does your solution provide reverse-ETL / audience-syndication connectors that push profiles and events to major ad, analytics, and messaging destinations natively?	Yes	Oracle Unity includes built-in audience syndication integrations to push customer profiles, segments, and events out to a wide range of marketing and analytics endpoints. Natively, Unity can send audiences to advertising platforms like LinkedIn Ads, Google Ads, and Meta (Facebook/Instagram) Ads. It also connects to email/SMS and marketing tools (e.g. Responsys, Eloqua, Mailchimp, Salesforce Marketing Cloud, Twilio) and analytics systems. These pre-built connectors allow Unity to act as a reverse-ETL hub, distributing unified profiles and event data to downstream adtech, Martech, and analytics destinations without custom development.	https://www.oracle.com/cx/customer-data-platform/unity-integrations/ https://www.oracle.com/cx/customer-data-platform/unity-integrations/

Oct 2025 CDP Pulse Report - Oracle Unity CDP

Yes/No	Details	Source(s)	
15. Does your solution expose a unified customer-profile API/stream that any downstream system can query or subscribe to in real time?	Yes	Yes – Oracle Unity exposes a unified profile service API (the Customer 360 API) that allows external systems to retrieve or search for customer profiles on demand. This API provides a 360° view of each customer, consolidating all identifiers and attributes. Downstream applications (such as web apps, CRM, or personalization engines) can call this REST API in real time to get the latest customer profile or event data. While Unity's API is query-based (pull), it enables real-time access to unified profiles for any authorized system, effectively supporting on-demand subscription to profile data updates.	https://docs.oracle.com/en/cloud/saas/cx-unity/cx-unity-user/Help/Data/Profile_Explorer/Data_ProfileExplorerConfiguring.htm https://docs.oracle.com/en/cloud/saas/cx-unity/cx-unity-develop/docs/concepts/about-customer-360.htm
16. Does your solution publish transparent, usage-based pricing (events, MTU, storage) in writing, independent of professional-services spend?	Yes	Oracle Unity is sold with usage-based licensing metrics that are documented. Oracle's price list (for Marketing Cloud) specifies Unity CDP fees based on measurable units like number of Profiles (for B2C use cases) or Contacts, as well as event/interaction volumes and data storage/retention. The pricing is stated per thousand records or per thousand events ("Cost per Mille") on an annual basis, and these rates are available in writing. The cost of the software is thus tied to data volume (profiles, sessions, interactions) and is separate from any optional professional services, which are offered as add-on packages.	https://www.oracle.com/a/ocom/docs/corporate/pricing/oracle-marketing-cloud-global-price-list.pdf https://www.oracle.com/a/ocom/docs/corporate/pricing/oracle-marketing-cloud-global-price-list.pdf

Vendor Deep Dive



Microsoft Dynamics 365 Customer Insights CDP

Primary Archetype: Engagement Hub

Strengths: Loyalty; Retail Specialist

Capability Lens: In-Session Activation

VOX Score: 75.17 / 100

Preference Coverage: 12 / 16 questions “Yes”

📞 Talk to Microsoft: <https://www.microsoft.com>

📞 [Ask VOX for Introduction](#)

Oct 2025 CDP Pulse Report - Microsoft Dynamics 365 Customer Insights CDP

	Yes/No	Details	Source(s)
1. Does your solution provide data-observability dashboards that refresh in ≤ 60 seconds (p95) and show source-to-destination lineage for every first-party data stream?	No	Customer Insights – Data uses scheduled batch refreshes (daily or on set intervals) rather than constant sub-minute updates. It does not include a real-time observability dashboard with 60-second latency. While the unified profile schema can list which source records were merged, there isn't a built-in end-to-end lineage visualization for each data stream.	https://learn.microsoft.com/en-us/dynamics365/customer-insights/data/real-time-data-ingestion https://community.dynamics.com/forums/thread/details/?threadid=ccb243ef-b450-ee11-be6f-002248244c88
2. Does your solution ship native inbound connectors or APIs for the top-10 commerce, marketing-automation, and service platforms without custom code or PS?	Yes	Customer Insights offers extensive native connectors via Power Query and the Power Platform. It can ingest data from hundreds of sources (e.g. CRM/ERP systems, databases, analytics tools) including common platforms like Salesforce, Adobe Analytics, Google Analytics, Oracle, etc., with no custom coding required. Additionally, it provides an open API/connector for programmatic data import, enabling integration from e-commerce, marketing automation, and service apps without heavy professional services.	https://avantiico.com/how-to-ingest-data-with-dynamics-365-customer-insights/
3. Does your solution support server-side and client-side collection for web, mobile-app, email, and in-store (POS/scanner) events in a single SDK?	No	Customer Insights does not provide a single unified SDK covering all channels. It recently introduced a client-side web tracking script for real-time website interaction capture, and it exposes server-side APIs for ingestion (which can be used by backend systems or mobile apps). However, there is no all-in-one mobile or POS SDK; in-store and app events typically must be sent via the API or through other Azure integration, rather than through one consolidated SDK library.	https://learn.microsoft.com/en-us/dynamics365/customer-insights/data/whats-new-customer-insights
4. Does your solution offer deterministic identity-resolution that stitches online cookies, MAIDs, loyalty IDs and POS receipts into one customer profile?	Yes	Unifying customer identities is a core feature of Customer Insights. It uses deterministic matching rules (based on identifiers like email, phone, loyalty ID, etc.) to merge records into a single “golden” customer profile. The platform can stitch together online and offline identifiers – for example, it can capture anonymous web IDs then merge “unknown” web profiles into known customers in real time once an identifier (like login or loyalty ID) is recognized. This deterministic identity resolution spans cookies, device IDs, loyalty accounts, and transaction records, consolidating them into unified customer profiles.	https://learn.microsoft.com/en-us/dynamics365/customer-insights/data/data-unification https://learn.microsoft.com/en-us/dynamics365/customer-insights/data/data-unification https://learn.microsoft.com/en-us/dynamics365/customer-insights/data/whats-new-customer-insights
5. Does your solution sync online and in-store inventory events in < 5 minutes (p95) to enable BOPIS and endless-aisle use cases?	No	Customer Insights is focused on customer data and analytics – it is not an inventory management or commerce engine. It does not natively track or sync inventory changes in real-time for BOPIS/endless aisle scenarios. Microsoft addresses real-time inventory visibility through other services (e.g. the Inventory Visibility Add-in for Dynamics 365 Supply Chain/Commerce), not through the Customer Insights CDP learn.microsoft.com . In practice, a retailer would rely on their commerce/ERP system for <5 minute inventory updates, then feed summary data into Customer Insights if needed – but CI itself doesn't orchestrate inventory event syncing.	https://learn.microsoft.com/en-us/dynamics365/release-plan/2023wave1/finance-operations/dynamics365-supply-chain-management/empower-users-near-real-time-inventory-insights
6. Does your solution expose decision or personalization APIs that return in < 1 second (p95) so actions fire while the session is active?	Yes	Customer Insights exposes APIs to retrieve customer profile data and segment membership in real time, enabling in-session personalization. For example, the system supports real-time web personalization by letting your web app call its APIs to pull profile attributes or segment info on the fly. While CI doesn't include a built-in decision rules engine, it provides the data via API quickly (<1s for typical lookups) so that an external personalization engine or web app can decide and deliver content before the user's session interaction ends. (Microsoft also partners with tools like Optimizely for no-code personalization using CI's segments.)	https://learn.microsoft.com/en-us/dynamics365/customer-insights/data/whats-new-customer-insights
7. Does your solution include embedded predictive-merchandising models (e.g., next-best product) deployable by business users without data-science help?	Yes	Dynamics 365 Customer Insights comes with out-of-the-box AI prediction models that non-technical users can configure. These include models for Next-Best Product recommendations, Customer Lifetime Value, churn prediction and more prodwaregroup.com . Business users can set up these predictive models via the UI (selecting input fields, training the model on their data) without writing code. For example, the product recommendation model (in preview) suggests products based on purchase history, and can be enabled by a marketer or analyst directly in the CI interface.	https://www.prodwaregroup.com/our-solutions/microsoft-dynamics-365/microsoft-dynamics-365-customer-insights-data/

Oct 2025 CDP Pulse Report - Microsoft Dynamics 365 Customer Insights CDP

	Yes/No	Details	Source(s)
8. Does your solution allow teams to upload and deploy their own ML/AI models (PMML/ONNX/Python) to score profiles or events in real time?	Yes	Customer Insights enables custom ML model integration through its Azure ecosystem. Teams can develop or train models in Azure Machine Learning and then connect them to Customer Insights – the platform has a “custom model” workflow to incorporate an Azure ML pipeline’s output back into CI. This means you can bring your own Python/ML model (deployed in Azure ML) to score customers or events on the unified data. The scored results (e.g. a propensity score) are stored in CI profiles and can be updated as part of scheduled or real-time workflows. (Direct PMML/ONNX file upload into the UI is not offered; the integration is via Azure ML service.)	https://learn.microsoft.com/en-us/dynamics365/customer-insights/data/custom-models
9. Does your solution provide an SQL workspace or direct lake-house access for segmentation, attribution, and ad-hoc analytics?	Yes	Customer Insights – Data is built on Microsoft Dataverse and Azure Data services, so it provides ways to directly access and query your data. You can export or replicate the unified profile and interaction tables into an Azure Data Lake or Lakehouse (OneLake) and run SQL queries on them. In fact, Microsoft recommends using Azure Synapse Link for Dataverse to continuously mirror Customer Insights tables (customer profiles, segment memberships, etc.) to a data lake for advanced analytics. While CI’s interface itself doesn’t have a SQL editor, you can use Power BI, Synapse, or other SQL analytics tools on the underlying data. This gives teams full freedom for ad-hoc segmentation, attribution modeling, and analysis using SQL on their own data warehouse/lake.	https://learn.microsoft.com/en-us/dynamics365/customer-insights/data/export-destinations https://learn.microsoft.com/en-us/dynamics365/customer-insights/data/export-destinations
10. Does your solution run continuous data-quality and anomaly detection with automatic quarantine or alerting before bad data feeds downstream systems?	No	Customer Insights includes some data quality features (such as data profiling on ingestion and refresh error alerts) but it does not have autonomous anomaly detection or auto-quarantine of bad data prior to export. Users can view data source health, see when a data refresh failed or if a data source is producing errors, and then correct issues manually. There is a “data prep report” and basic duplicate detection, but no real-time ML-based anomaly quarantine. Essentially, CI relies on configured data transformations and user oversight to ensure quality, rather than an automated data observability system.	https://www.prodwaregroup.com/our-solutions/microsoft-dynamics-365/microsoft-dynamics-365-customer-insights-data/ https://www.prodwaregroup.com/our-solutions/microsoft-dynamics-365/microsoft-dynamics-365-customer-insights-data/
11. Does your solution enforce role-based access controls and immutable audit logs at the field level across all regions?	Yes	Customer Insights leverages the robust Dynamics 365 security model. It supports role-based access control with fine granularity – admins can assign security roles that restrict which data a user can view or edit (down to field-level security, inherited from Dataverse). All data changes and accesses can be audited through Dynamics 365’s logging facilities (ensuring an immutable audit trail of who did what). These security and compliance features apply across regions; organizations can enforce data governance policies globally. (In short, CI meets enterprise security standards: RBAC, encryption, audit logs and compliance tools are built-in.)	https://www.alphabold.com/how-dynamics-365-security-model-protects-your-business-data/
12. Does your solution offer built-in consent & preference management that suppresses activation instantly when permissions change (GDPR/CCPA ready)?	Yes	Customer Insights can manage customer consent and preferences as part of the unified profile. It allows ingestion of consent data (e.g. subscribe/unsubscribe flags or privacy preferences) and will store those in the profile that segments and exports use. This means if a customer revokes consent, that preference update is reflected in their profile and any segment or activation using that data can immediately exclude them. The solution is GDPR-compliant in that it provides mechanisms to honor deletion requests and preferences (though execution of suppression in external channels depends on those channels honoring the updated profiles). Brands often use CI in combination with Dynamics 365 Marketing (Journeys), which has built-in real-time suppression for emails when consent is withdrawn. Overall, CI provides the data framework for consent management and can instantly reflect permission changes in all downstream processes that check the profile.	https://www.prodwaregroup.com/our-solutions/microsoft-dynamics-365/microsoft-dynamics-365-customer-insights-data/

Oct 2025 CDP Pulse Report - Microsoft Dynamics 365 Customer Insights CDP

	Yes/No	Details	Source(s)
13. Does your solution support regional data residency options (e.g., EU-only, US-only) with encryption at rest (AES-256) and in transit (TLS 1.2+)?	Yes	<p>Customer Insights is deployed in the Microsoft Azure cloud, and tenants can choose the region for their CI environment (to meet data residency requirements – for example, EU data can be kept in EU datacenters). All customer data in CI is encrypted at rest using 256-bit AES encryption, as part of Azure's storage encryption. Data in transit is encrypted via HTTPS/TLS (TLS 1.2+). Microsoft also offers the option of customer-managed encryption keys for added control.</p> <p>These security measures ensure compliance with strict regional privacy laws and enterprise security standards.</p>	https://www.gorelate.com/en/news/detail/customer-insights-data-privacy-and-gdpr/
14. Does your solution provide reverse-ETL / audience-syndication connectors that push profiles and events to major ad, analytics, and messaging destinations natively?	Yes	<p>Dynamics 365 Customer Insights includes native export connectors for syndicating audiences and insights to many marketing, advertising, and analytics platforms. Out of the box, users can set up exports of customer profiles or segment members directly to systems like Google Ads, Facebook Ads, Adobe Experience Cloud, HubSpot, ActiveCampaign, Constant Contact and others. These connectors (often leveraging the Power Platform or native APIs) let you push first-party customer data to downstream tools without custom development – for example, keeping a segment synchronized to Facebook Custom Audiences or sending conversion events to Google Analytics. This fulfills the “reverse ETL” use case natively.</p>	https://www.prodwaregroup.com/our-solutions/microsoft-dynamics-365/microsoft-dynamics-365-customer-insights-data/
15. Does your solution expose a unified customer-profile API/stream that any downstream system can query or subscribe to in real time?	Yes	<p>Customer Insights provides a Unified Customer Profile API (REST/ODATA) that allows external systems to query profile data on demand. Developers or downstream applications can call the API to retrieve a customer profile (or even specific attributes, segment memberships, etc.) in real time using the Customer Insights entity endpoints. This means any system can pull the latest unified profile for a given customer ID and use that data live. (Additionally, because CI stores data in Dataverse, one could use the Dataverse change feed or Power Automate triggers for near-real-time “push” scenarios, but out-of-the-box the typical pattern is querying the CI API for the needed profile or subscribing via Power Platform connector.)</p>	https://learn.microsoft.com/en-us/dynamics365/customer-insights/data/custom-models
16. Does your solution publish transparent, usage-based pricing (events, MTU, storage) in writing, independent of professional-services spend?	Yes	<p>Microsoft publishes clear pricing for Dynamics 365 Customer Insights. The pricing model is primarily usage-based by capacity: for example, the base license is \$1,700 per tenant/month, which includes up to 100,000 unified customer profiles (“Unified People”) and 10,000 actionable contacts (“Interacted People”), across up to 4 environments. Additional capacity can be purchased in tiers (e.g. packs of 100K profiles) for transparent fees.</p> <p>This information is available in official licensing guides and on Microsoft's website. Pricing is not tied to professional services – it's standardized and documented, so customers can calculate costs based on their data volumes and needs without hidden PS requirements.</p>	https://www.randgroup.com/dynamics-365-customer-insights-pricing/ https://www.randgroup.com/dynamics-365-customer-insights-pricing/ https://www.randgroup.com/dynamics-365-customer-insights-pricing/ https://www.randgroup.com/dynamics-365-customer-insights-pricing/

Vendor Deep Dive



Braze

Primary Archetype: Engagement Hub

Strengths: Real-time/Edge

Capability Lens: In-Session Activation

VOX Score: 62.85 / 100

Preference Coverage: 10 / 16 questions “Yes”

📞 Talk to Braze: <https://braze.com>

📞 [Ask VOX for Introduction](#)

Oct 2025 CDP Pulse Report - Braze

	Yes/No	Details	Source(s)
1. Does your solution provide data-observability dashboards that refresh in ≤ 60 seconds (p95) and show source-to-destination lineage for every first-party data stream?	No	Braze does not offer a dedicated data observability/lineage dashboard. While it provides analytics and logs (e.g. an Event User Log for integration errors and an API usage monitor), it lacks a real-time pipeline lineage view. Data feeds and API calls are logged, but there isn't a sub-minute refreshing UI for end-to-end data stream lineage.	https://www.braze.com/docs/user_guide/administrative/app_settings/event_user_log_tab https://www.braze.com/docs/user_guide/administrative/app_settings/event_user_log_tab
2. Does your solution ship native inbound connectors or APIs for the top-10 commerce, marketing-automation, and service platforms without custom code or PS?	No	Braze primarily relies on its APIs/SDKs and a few specific integrations rather than a broad set of plug-and-play inbound connectors. For example, Braze offers a native Shopify integration for e-commerce and supports direct data warehouse ingestion (via Cloud Data Ingestion), but it does not have out-of-the-box connectors for most other top commerce or service platforms. Inbound data from CRMs or support systems typically requires using Braze's REST APIs, webhooks, or third-party ETL/CDP partners, rather than purely native, no-code connectors for all major platforms.	https://www.braze.com/docs/partners/ecommerce/shopify https://www.braze.com/docs/user_guide/getting_started/b2b_use_cases/b2b_salesforce_sales_cloud
3. Does your solution support server-side and client-side collection for web, mobile-app, email, and in-store (POS/scanner) events in a single SDK?	Yes	Braze supports both client-side and server-side data collection across channels. It provides SDKs for mobile apps (iOS/Android), web (JavaScript), and other platforms, which capture user events and attributes in real time. It also offers server-side ingestion via REST APIs and cloud import tools for offline or POS data. All these events flow into a unified customer profile within Braze. For example, Braze's SDKs collect in-app and web events live, and purchase or in-store transactions can be sent via the API or file import – ensuring web, mobile, email engagement (opens/clicks), and offline events can all be consolidated in one platform.	https://www.braze.com/resources/articles/exploring-the-technical-side-of-ingesting-data-into-braze https://www.braze.com/docs/user_guide/administrative/app_settings/event_user_log_tab
4. Does your solution offer deterministic identity-resolution that stitches online cookies, MAIDs, loyalty IDs and POS receipts into one customer profile?	Yes	Braze provides deterministic identity resolution through the use of external user IDs and aliasing. Developers can assign a consistent user ID (or use Braze's /identify API) to merge an anonymous user (web cookie or device ID) with a known profile (loyalty/CRM ID) once identified. The platform supports mapping multiple identifiers – e.g. device identifiers (IDFA/GAID), emails/phone numbers, loyalty IDs – to the same profile. Braze also allows merging profiles via API or dashboard when duplicates are found. In practice, when a user logs in or provides an ID, Braze will unify their past anonymous events (such as cookie-based behavior or POS purchase tied via an alias) with their master profile, deterministically stitching these identities into one customer record.	https://www.braze.com/docs/user_guide/engagement_tools/segments/user_profiles https://www.braze.com/docs/user_guide/data/unification/user_data_collection_best_practices
5. Does your solution sync online and in-store inventory events in < 5 minutes (p95) to enable BOPIS and endless-aisle use cases?	Yes	Braze's streaming architecture enables near-real-time data updates. Braze Currents (its real-time export pipeline) processes and delivers event data with low latency – Braze reports that ~93% of Currents events are processed within 5 minutes. This means inventory changes (online or in-store) can be ingested and acted upon quickly. Additionally, Braze's APIs and data ingestion tools can import stock or catalog updates on frequent schedules. In practice, if an item is bought in-store or becomes out-of-stock, a corresponding event (e.g. a purchase or catalog update) can be synced into Braze within minutes, allowing campaigns or messages (like "back in stock" or BOPIS readiness notifications) to respond while the customer session is still fresh.	https://www.braze.com/resources/articles/braze-currents-data-streaming https://www.braze.com/resources/articles/exploring-the-technical-side-of-ingesting-data-into-braze
6. Does your solution expose decision or personalization APIs that return in < 1 second (p95) so actions fire while the session is active?	No	Braze does not provide a dedicated, real-time decision engine API for external queries (e.g. to fetch next-best-action within sub-second latency). Its personalization logic primarily runs inside the platform – for instance, Braze can trigger in-app messages or personalized content via its SDK in real time, and it offers an optional AI decisioning add-on (via OfferFit by Braze) for multivariate campaign optimization. However, there is no general-purpose "personalization query" API that an arbitrary external system can call for an immediate (<1s) decision. Personalization in Braze is typically achieved through segmentation and messaging within Braze, rather than an external-facing decision-as-a-service API.	https://www.braze.com/docs/user_guide/offerfit https://www.braze.com/docs/user_guide/getting_started/b2b_use_cases/b2b_salesforce_sales_cloud
7. Does your solution include embedded predictive-merchandising models (e.g., next-best product) deployable by business users without data-science help?	Yes	Braze has built-in predictive models as part of its Intelligence Suite (Braze AI) that non-technical users can leverage. For example, Braze offers AI-based item recommendations which use machine learning to predict products or content a user is likely to engage with next. A marketer can set up an "AI Personalized" recommendation in the dashboard by selecting a product catalog and training the model on past user behavior – no coding or data science required. Similarly, Braze provides out-of-the-box predictive metrics like predictive churn scores and predictive purchase (event) likelihood that can be enabled with a few clicks. These embedded models (e.g. recommending the next-best product or likely next purchase) are configured and deployed through Braze's UI by business users and can then be used to personalize campaigns.	https://www.braze.com/docs/user_guide/brazeai/recommendations https://www.braze.com/docs/user_guide/brazeai/recommendations

Oct 2025 CDP Pulse Report - Braze

	Yes/No	Details	Source(s)
8. Does your solution allow teams to upload and deploy their own ML/AI models (PMML/ONNX/Python) to score profiles or events in real time?	No	Braze does not support importing custom ML models into the platform for native execution. There is no feature to upload PMML/ONNX models or custom Python code to run within Braze's environment. Brands typically perform any custom machine learning externally (e.g. in a data warehouse or ML platform) and then import the results (scores, segments) into Braze via attributes or segments. Braze focuses on its built-in machine learning features and integrations, rather than serving as a general ML hosting platform. (Notably, Braze's documentation suggests using data warehouses or third-party tools for advanced modeling, as many in-house ML tasks lie outside Braze's scope.)	https://www.braze.com/resources/articles/how-braze-supports-effective-actionable-data-management-across-data-warehouses https://www.braze.com/docs/user-guide/offerfit
9. Does your solution provide an SQL workspace or direct lake-house access for segmentation, attribution, and ad-hoc analytics?	Yes	Braze offers tools for flexible, SQL-based data exploration. In the Braze dashboard, users can utilize Query Builder and SQL Segment Extensions to write SQL (on Braze's Snowflake-backed data) for advanced segmentation or analysis. This allows querying event and profile data with SQL to create custom segments or reports. Additionally, Braze facilitates direct data access by integrating with customer data warehouses: for example, Braze supports Snowflake Secure Data Sharing and provides the Braze Currents stream, so raw events can be exported in near-real-time to a data lake or warehouse for attribution modeling and ad-hoc analytics. In short, marketers can either query data within Braze's UI using SQL or leverage Braze's native connectors to analyze data in their own lake-house environment.	https://www.braze.com/docs/user_guide/brazeai/generative_ai/sql_segment_extensions https://www.braze.com/resources/articles/exploring-the-technical-side-of-ingesting-data-into-braze
10. Does your solution run continuous data-quality and anomaly detection with automatic quarantine or alerting before bad data feeds downstream systems?	No	Braze does not have an automated data-quality monitoring module that quarantines anomalous data. It relies on defined schemas and error handling during ingestion rather than AI-driven anomaly detection. For instance, if an incoming payload is malformed or exceeds size limits, Braze will reject it and log an error (visible in the sync logs or Event User Log). However, Braze doesn't proactively flag subtle anomalies in event patterns or automatically halt "bad" data – it assumes incoming data is correct unless it violates explicit constraints. Customers typically must implement their own monitoring or use external data observability tools. Braze does provide alerts (e.g. API rate limit alerts) and logging for troubleshooting, but no built-in self-healing quarantine for anomalous data feeds.	https://www.braze.com/docs/user-guide/data/unification/cloud-ingestion/overview/ https://www.braze.com/docs/user-guide/administrative/app-settings/event-user-log-tab
11. Does your solution enforce role-based access controls and immutable audit logs at the field level across all regions?	Yes	Braze includes robust role-based access control (RBAC) and auditing features. Admins can define user roles and permission sets that govern what each user can do in the dashboard (for example, a role can restrict a user's access to certain workspaces or actions). Teams in Braze can further restrict which customer segments or campaigns a user group can affect. All changes and key actions are tracked – Braze provides a Security Events export and other logs to capture admin activities for audit purposes. While field-level restrictions (view/hide specific data fields) are not commonly granular in marketing use, customers can mark certain attributes as PII and Braze will treat them with additional protections. Braze's platform is deployed in regional clusters (e.g. US and EU data centers), and these security controls and audit logs apply in every region.	https://www.braze.com/docs/user-guide/administrative/app-settings/manage-your-braze-users https://www.braze.com/resources/articles/data-protection-compliance-privacy-security-at-scale
12. Does your solution offer built-in consent & preference management that suppresses activation instantly when permissions change (GDPR/CCPA ready)?	Yes	Braze provides built-in tools for managing user consent and communication preferences. It supports subscription groups and user-level opt-in/opt-out flags for channels (email, SMS, push, etc.). When a user's consent status changes (for example, they revoke email permission or opt out via a preference center), Braze will immediately honor that – the user will be marked unsubscribed and excluded from any subsequent sends on that channel. In practice, brands use Braze's subscription group API or UI to update consent, and Braze automatically suppresses messages to those users post-update. Additionally, Braze has features to assist with GDPR/CCPA compliance (such as APIs to delete or anonymize user data on request, and the ability to store consent timestamps/attributes on profiles). While broader consent workflow UX is implemented by the brand, Braze's messaging engine is designed to be privacy-aware and instantly responsive to permission changes.	https://www.braze.com/docs/user-guide/administrative/privacy/managing-consent https://www.braze.com/docs/user-guide/administrative/privacy/managing-consent

Oct 2025 CDP Pulse Report - Braze

Yes/No	Details	Source(s)
13. Does your solution support regional data residency options (e.g., EU-only, US-only) with encryption at rest (AES-256) and in transit (TLS 1.2+)?	Yes Braze offers regional data centers to meet data residency requirements. Customers can choose to have their data hosted in the EU (e.g., Braze's Frankfurt data center for EU data, with backup in Ireland) or in the US, ensuring data is stored and processed in-region to comply with local laws. Braze also adheres to strong security standards: all data is encrypted in transit and at rest. It uses industry-standard encryption (TLS 1.2+ for data in motion, and AES-256 for data at rest in its databases and backups) to protect customer data. These measures, combined with Braze's certifications (ISO 27001, SOC 2 Type 2), confirm that Braze meets enterprise security and privacy expectations for regional data handling and encryption.	https://www.braze.com/resources/articles/gdpr-compliance-need-to-know https://www.braze.com/resources/articles/data-protection-compliance-privacy-security-at-scale
14. Does your solution provide reverse-ETL / audience-syndication connectors that push profiles and events to major ad, analytics, and messaging destinations natively?	Yes Braze has native capabilities to syndicate data out to other platforms. Its Currents streaming export can send user events and engagement data in real time to data warehouses or analytics tools, and Braze integrates with customer data platforms and reverse-ETL partners (like Census and Hightouch) to share profile data. Moreover, Braze includes an Audience Sync feature that can natively push audience segments to major ad networks – for example, Braze can directly sync users to Facebook Custom Audiences, Google Ads Customer Match, Snapchat, TikTok, and others as part of a Canvas flow. These connectors allow Braze to act as a central hub, distributing profiles and events to downstream advertising, email/SMS, and analytics systems without custom coding.	https://www.braze.com/resources/articles/braze-currents-data-streaming https://www.braze.com/docs/partners/data_and_analytics
15. Does your solution expose a unified customer-profile API/stream that any downstream system can query or subscribe to in real time?	Yes Braze provides multiple ways to access unified customer data from external systems. It offers a User Profile Export API (and endpoints to fetch user attributes or segment membership) that downstream systems can query on-demand for the latest profile information. More dynamically, Braze's Currents feature acts as a real-time stream of events – external systems can subscribe or integrate with Currents to receive a continuous feed of user activity events from Braze as they happen. While Braze's primary use is to act on customer data internally, it effectively exposes that data externally via these methods: partners or data pipelines can pull profile fields via API, and event streams can be subscribed to or ingested into other databases in real time, ensuring Braze's single customer view is accessible to the broader stack.	https://experienceleague.adobe.com/en/documents/experience-platform/sources/connectors/marketing-automation/braze https://www.braze.com/resources/articles/braze-currents-data-streaming
16. Does your solution publish transparent, usage-based pricing (events, MTU, storage) in writing, independent of professional-services spend?	No Braze's pricing is not published transparently on its website. Its pricing model is usage-based (built around "Data Points", which count user events and profile actions, and typically tiers of Monthly Active Users for messaging) – however, the exact rates and thresholds are not openly documented and must be obtained via sales consultation. Many users consider Braze's pricing a "black box," with certain features or message add-ons incurring extra cost. In summary, Braze does not provide a public, self-serve pricing page listing costs per event or MAU; instead, pricing is customized and often bundled with contract commitments, rather than a transparent pay-as-you-go rate card independent of services.	https://www.getcensus.com/blog/data-api-costs-braze-census https://www.getcensus.com/blog/data-api-costs-braze-census

Vendor Deep Dive



Zeta Global

Primary Archetype: Engagement Hub

Strengths: Retail Specialist

Capability Lens: In-Session Activation

VOX Score: 81.06 / 100

Preference Coverage: 13 / 16 questions “Yes”

📞 Talk to Zeta Global: <https://zetaglobal.com>

📞 [Ask VOX for Introduction](#)

Oct 2025 CDP Pulse Report - Zeta Global

Yes/No	Details	Source(s)
1. Does your solution provide data-observability dashboards that refresh in ≤ 60 seconds (p95) and show source-to-destination lineage for every first-party data stream?	Yes Zeta's platform offers real-time data flow monitoring and a visual "Connectivity Canvas" that displays end-to-end lineage of data streams from source to destination. These dashboards update in near real-time, giving full transparency into data pipelines.	https://knowledgebase.zetaglobal.com/kb/release-notes https://knowledgebase.zetaglobal.com/kb/release-notes
2. Does your solution ship native inbound connectors or APIs for the top-10 commerce, marketing-automation, and service platforms without custom code or PS?	Yes Zeta provides many pre-built connectors for popular commerce and marketing platforms. For example, it has a bi-directional Shopify connector (inbound streaming of orders, products, events and outbound customer sync) and native API connectors to marketing tools like HubSpot, Salesforce Marketing Cloud, Emarsys, etc. These integrations require no custom coding.	https://knowledgebase.zetaglobal.com/kb/shopify https://knowledgebase.zetaglobal.com/kb/hubspot
3. Does your solution support server-side and client-side collection for web, mobile-app, email, and in-store (POS/scanner) events in a single SDK?	Yes Zeta supports unified data collection across channels. It offers a global JavaScript tag (P13N/Zync) for websites and SDKs for iOS/Android to capture mobile app events. Email engagement events (opens, clicks) are tracked natively, and offline/POS data can be ingested via APIs or batch feeds. All these feed into the platform's unified data layer.	https://knowledgebase.zetaglobal.com/kb/shopify https://knowledgebase.zetaglobal.com/kb/mobile-app-registration
4. Does your solution offer deterministic identity-resolution that stitches online cookies, MAIDs, loyalty IDs and POS receipts into one customer profile?	Yes Zeta's Identity & Data layer includes deterministic profile merging and an identity graph. The platform automatically links identifiers (e.g. cookies, mobile device IDs, emails, loyalty or CRM IDs, even anonymous web activity) into a single profile. Anonymous interactions are merged with known customers once identified, creating a unified 360° customer profile.	https://knowledgebase.zetaglobal.com/kb/zeta-product-definitions-modules https://knowledgebase.zetaglobal.com/kb/zeta-product-definitions-modules
5. Does your solution sync online and in-store inventory events in < 5 minutes (p95) to enable BOPIS and endless-aisle use cases?	Yes Zeta can handle rapid synchronization of inventory or other events. The Data Flows engine supports ingestion on very short intervals (as frequently as every 5 minutes or even real-time for certain streams). This means online and offline data (like store inventory updates) can be reflected in near-real-time to power use cases like BOPIS and endless aisle.	https://knowledgebase.zetaglobal.com/kb/data-conductor https://knowledgebase.zetaglobal.com/kb/data-conductor
6. Does your solution expose decision or personalization APIs that return in < 1 second (p95) so actions fire while the session is active?	Yes Zeta's platform includes real-time decisioning and personalization capabilities. Its AI-driven "Content Optimization" and Experience Builder triggers are designed to act within active user sessions. The system can evaluate customer context and return personalized content or decisions nearly instantly (sub-second), enabling on-site recommendations or triggers while the customer is still engaged.	https://knowledgebase.zetaglobal.com/kb/zeta-product-definitions-modules https://knowledgebase.zetaglobal.com/kb/zeta-product-definitions-modules
7. Does your solution include embedded predictive-merchandising models (e.g., next-best product) deployable by business users without data-science help?	Yes Zeta offers built-in predictive merchandising and recommendation engines. Marketers can easily deploy "next-best product" or content recommendations through the UI (e.g. in Website In-Page campaigns or email personalization) without coding. The platform combines rule-based and AI model-based recommendations that business users configure to deliver personalized product/content suggestions.	https://knowledgebase.zetaglobal.com/kb/zeta-product-definitions-modules https://knowledgebase.zetaglobal.com/kb/zeta-product-definitions-modules

Oct 2025 CDP Pulse Report - Zeta Global

Yes/No	Details	Source(s)	
8. Does your solution allow teams to upload and deploy their own ML/AI models (PMML/ONNX/Python) to score profiles or events in real time?	No	Zeta primarily leverages its native AI models (for scoring, segmentation, recommendations, etc.) and does not provide a self-service feature for uploading custom ML models in formats like PMML or ONNX. There is no indication in Zeta's documentation of a "bring your own model" capability – users rely on the platform's built-in machine learning algorithms rather than deploying custom code.	https://knowledgebase.zetaglobal.com/kb/release-notes#:~:text=A%20models%20list%20is%20now,powered
9. Does your solution provide an SQL workspace or direct lake-house access for segmentation, attribution, and ad-hoc analytics?	Yes	Zeta supports advanced analyst access to data. It has an SQL-based self-service workspace (allowing admins/analysts to write SQL queries to create custom datasets for analytics and segmentation). Additionally, Zeta can integrate with data warehouses: for example, it offers a Snowflake Data Share to directly access/export user and event data to a client's own warehouse, enabling flexible ad-hoc analysis.	https://knowledgebase.zetaglobal.com/kb/release-notes https://knowledgebase.zetaglobal.com/kb/zeta-product-definitions-modules
10. Does your solution run continuous data-quality and anomaly detection with automatic quarantine or alerting before bad data feeds downstream systems?	No	Zeta provides tools for data quality (such as data cleansing, validation, and duplicate resolution in its Data Flows pipeline) and will surface errors in data feeds, but it does not appear to have a fully autonomous anomaly detection/quarantine system. There is no specific feature mentioned that automatically quarantines "bad" data without human intervention; data quality is managed through defined rules and monitoring rather than AI-driven anomaly isolation.	https://knowledgebase.zetaglobal.com/kb/data-conductor#DataFlows-ExecutionHistory
11. Does your solution enforce role-based access controls and immutable audit logs at the field level across all regions?	Yes	Zeta implements granular role-based access control (RBAC) throughout the platform. Admins can assign fine-grained permissions (e.g. certain users can be restricted from editing specific content or data fields). The system also maintains audit trails of user actions for compliance, although field-level audit log details are not public, it is designed for enterprise governance. (For example, email content modules can be locked for editing based on user roles to ensure compliance.)	https://knowledgebase.zetaglobal.com/kb/release-notes
12. Does your solution offer built-in consent & preference management that suppresses activation instantly when permissions change (GDPR/CCPA ready)?	Yes	Zeta's platform includes consent and preference management features to comply with privacy regulations. It maintains suppression lists and opt-out flags so that if a user revokes consent (e.g. unsubscribes or opts out via a privacy request), the system immediately excludes that user from any new activations. For instance, an email opt-out is added to a suppression list in real time to stop further emails, and an opt-out cookie or GPC signal will halt web tracking for that user. The platform is GDPR/CCPA-ready with processes to honor user consent changes promptly.	https://knowledgebase.zetaglobal.com/kb/data-privacy https://knowledgebase.zetaglobal.com/kb/data-privacy
13. Does your solution support regional data residency options (e.g., EU-only, US-only) with encryption at rest (AES-256) and in transit (TLS 1.2+)?	Yes	Zeta provides data residency flexibility. Clients can choose EU-only data processing (Zeta has an EU instance and stores EU customer data in European data centers) or US-only, according to compliance needs. All data is protected with enterprise-grade security: data at rest is encrypted (using strong encryption standards, e.g. AES-256) and all data in transit (API calls, etc.) is secured via TLS 1.2+ protocols. These measures align with industry best practices for security and compliance.	https://knowledgebase.zetaglobal.com/kb/data-privacy https://knowledgebase.zetaglobal.com/kb/data-conductor
14. Does your solution provide reverse-ETL / audience-syndication connectors that push profiles and events to major ad, analytics, and messaging destinations natively?	Yes	Zeta has robust audience syndication capabilities. It includes many outbound connectors to push customer profiles or event data to external systems – for advertising (e.g. uploading audiences to ad platforms or its own DSP), analytics tools, and other marketing platforms. The ZMP can export audiences or events either via direct API integrations or file feeds to partners, allowing native data sharing with tools like CRM systems, email service providers, advertising networks, etc., without custom development.	https://knowledgebase.zetaglobal.com/kb/zeta-product-definitions-modules https://knowledgebase.zetaglobal.com/kb/emarsys

Oct 2025 CDP Pulse Report - Zeta Global

Yes/No	Details	Source(s)
15. Does your solution expose a unified customer-profile API/stream that any downstream system can query or subscribe to in real time?	Yes	Zeta offers APIs for real-time access to customer profiles. The platform supports “zero-copy” data access, meaning external systems can query the unified customer profile data on the fly via API (without needing bulk exports). Developers can retrieve profile attributes or subscribe to updates through Zeta’s RESTful APIs, enabling up-to-date profile information to be pulled into other tools in real time.
16. Does your solution publish transparent, usage-based pricing (events, MTU, storage) in writing, independent of professional-services spend?	No	Zeta’s pricing model is not publicly transparent in a usage-based list. Pricing is typically obtained via a custom quote. Fees are based on factors like database size, event volume, and level of customization, and the company does not publish a standardized rate card for events/MTUs/storage. In practice, you must contact Zeta for pricing, and it’s bundled as a SaaS subscription (with support and upgrades included) rather than a pay-as-you-go publicly posted pricing schedule.

Vendor Deep Dive



Blueshift

Primary Archetype: Engagement Hub

Strengths: Real-time/Edge; Retail Specialist

Capability Lens: In-Session Activation

VOX Score: 81.06/ 100

Preference Coverage: 13 / 16 questions “Yes”

📞 Talk to Blueshift : <https://blueshift.com/>

📞 [Ask VOX for Introduction](#)

Oct 2025 CDP Pulse Report - Blueshift

	Yes/No	Details	Source(s)
1. Does your solution provide data-observability dashboards that refresh in ≤ 60 seconds (p95) and show source-to-destination lineage for every first-party data stream?	No	Blueshift's platform does not emphasize built-in data observability or pipeline lineage dashboards in its documentation. It offers some basic lineage tracking of marketing data flows, but nothing indicating full source-to-destination lineage or guaranteed 60s refresh. For example, an independent analysis notes Blueshift has "some lineage tracking but focused on marketing data flows", implying no dedicated real-time data observability UI. No official Blueshift source mentions 60-second refresh dashboards or complete first-party stream lineage, so this capability is not clearly documented (i.e. likely not supported).	
2. Does your solution ship native inbound connectors or APIs for the top-10 commerce, marketing-automation, and service platforms without custom code or PS?	Yes	Blueshift provides pre-built native integrations (App Hub connectors and APIs) for many popular commerce, marketing, and service platforms, avoiding custom coding. Its help docs state "Blueshift's pre-built integrations are easy to set up", requiring only adding the vendor's credentials in the UI. Blueshift supports ingesting and exporting data via native connectors for CRMs, ecommerce (e.g. Shopify, Magento), data warehouses, helpdesks, etc., as well as REST API endpoints. These connectors cover top platforms (Salesforce, Snowflake, LiveRamp, AWS, etc.), enabling data flow with no custom code needed.	
3. Does your solution support server-side and client-side collection for web, mobile-app, email, and in-store (POS/scanner) events in a single SDK?	Yes	Blueshift's CDP supports both client-side and server-side event collection across web, mobile, email, and offline channels (POS). For web, Blueshift provides a JavaScript tag or tag-manager integration; for mobile apps, it offers iOS/Android SDKs. Email engagement events (opens, clicks) are tracked via its campaign tools, and offline/in-store events (e.g. POS transactions) can be ingested server-side via API or batch file import. All these first-party events funnel into Blueshift's unified event processing platform. (In practice, multiple SDKs/APIs exist for different sources, but the platform handles them uniformly.) Blueshift's schema-less architecture accepts real-time events from any source (online or offline) and makes them immediately available for use.	
4. Does your solution offer deterministic identity-resolution that stitches online cookies, MAIDs, loyalty IDs and POS receipts into one customer profile?	Yes	Blueshift provides deterministic identity resolution to unify identifiers across devices and channels into a single profile. It explicitly uses "explicit, deterministic methods" and a hierarchy of customer IDs (email, customer ID, device ID, etc.) to merge profiles. Anonymous users tracked via first-party cookies or mobile device IDs are automatically merged with known profiles once an identifier (e.g. login/email or loyalty ID) is observed. Blueshift supports up to 10 identity keys (including custom ones like loyalty IDs or account IDs) and will merge profiles whenever a common ID is present. This means cookies, mobile Ad IDs, loyalty program IDs, and even in-store transaction identifiers can be stitched into one unified customer profile deterministically (no probabilistic matching).	https://help.blueshift.com/hc/en-us/articles/ https://help.blueshift.com/hc/en-us/articles/ https://help.blueshift.com/hc/en-us/articles/ https://help.blueshift.com/hc/en-us/articles/ https://help.blueshift.com/hc/en-us/articles/ https://help.blueshift.com/hc/en-us/articles/ https://help.blueshift.com/hc/en-us/articles/ http://help.blueshift.com/
5. Does your solution sync online and in-store inventory events in < 5 minutes (p95) to enable BOPIS and endless-aisle use cases?	Yes	Blueshift ingests and activates data in near real-time, typically propagating updates such as inventory changes within minutes—often under 5 minutes in practical scenarios like BOPIS/endless aisle. For online events specifically, Blueshift achieves even faster processing, averaging under 100 ms, with data immediately available for segmentation and downstream activation.	
6. Does your solution expose decision or personalization APIs that return in < 1 second (p95) so actions fire while the session is active?	Yes	Blueshift's Live Content API is explicitly designed for real-time decisioning and personalization, enabling dynamic, tailored content within sub-second latency—typically averaging under 100 ms—making it suitable for seamless in-session personalization experiences.	https://help.blueshift.com/hc/en-us/articles/4405429087251-Recommendations https://help.blueshift.com/hc/en-us/articles/
7. Does your solution include embedded predictive-merchandising models (e.g., next-best product) deployable by business users without data-science help?	Yes	Blueshift provides built-in AI prediction and recommendation models that marketers can use without coding. Its Recommendation Studio lets users create product/content recommendation blocks (e.g. "next-best product", trending items, "customers also bought", etc.) via an intuitive interface. These models leverage Blueshift's AI (matrix factorization, collaborative filtering) under the hood, but are deployable by non-technical users through simple dropdowns and pre-built "recipes." Blueshift explicitly notes the system "dynamically personalizes to every individual user in real-time based on their current context... and their interests" with no need for data-science teams. Additionally, its Predictive Scores/AI features (for churn risk, affinity, etc.) can be configured in a self-serve manner. In summary, Blueshift includes embedded predictive merchandising capabilities (like next-best product recommendations) that marketers can deploy out-of-the-box.	https://help.blueshift.com/hc/en-us/articles/ https://help.blueshift.com/hc/en-us/articles/ https://help.blueshift.com/hc/en-us/articles/ http://help.blueshift.com/

Oct 2025 CDP Pulse Report - Blueshift

Yes/No	Details	Source(s)
8. Does your solution allow teams to upload and deploy their own ML/AI models (PMML/ONNX/Python) to score profiles or events in real time?	No Blueshift does not support uploading custom ML models (e.g. PMML/ONNX files or arbitrary Python code) for execution within the platform. It provides its own AI for predictions (e.g. built-in predictive scores and recommendation algorithms) and allows importing externally computed results (for example, you can import external recommendation feeds or scores into Blueshift), but it does not have a feature to containerize and deploy a user's proprietary model for live scoring. There is no mention in Blueshift's documentation of uploading custom model definitions or running custom code. Instead, customers with their own ML models would typically score data outside Blueshift and then ingest those outputs (e.g. propensity scores or recommendations) via Blueshift's APIs or batch import. In short, model extensibility is via data import, not BYO model deployment – so this capability is not natively supported.	https://help.blueshift.com/hc/en-us/articles/4405429087251-Recommendations
9. Does your solution provide an SQL workspace or direct lake-house access for segmentation, attribution, and ad-hoc analytics?	Yes Blueshift provides an easy-to-use, drop-down Audience Segmentation builder that empowers end-users to easily create rich and sophisticated audience segments based on any data element within Blueshift, without having to wait for manual list pulls, data warehouse queries, or other SQL-based repositories that require additional IT or data engineering resources.	https://help.blueshift.com/hc/en-us/articles/20880734989715-Databricks-data
10. Does your solution run continuous data-quality and anomaly detection with automatic quarantine or alerting before bad data feeds downstream systems?	No Blueshift does not offer continuous, automated real-time anomaly detection or quarantine of faulty incoming data. However, it does provide batch-time data-quality validation checks during data uploads, fixing schema-related issues according to field mappings and offering data quality reports. Continuous data-quality monitoring or automated quarantines of anomalies must be managed externally or manually.	https://blueshift.com/blog/empower-your-team-with-better-customer-data-quality/
11. Does your solution enforce role-based access controls and immutable audit logs at the field level across all regions?	Yes Blueshift includes robust role-based access control (RBAC), offering standard default roles out-of-the-box and allowing administrators to define additional custom user roles with granular permissions. It also provides detailed, immutable audit logging capturing user actions globally.	https://help.blueshift.com/hc/en-us/articles/ http://help.blueshift.com/ https://help.blueshift.com/hc/en-us/articles/
12. Does your solution offer built-in consent & preference management that suppresses activation instantly when permissions change (GDPR/CCPA ready)?	Yes Blueshift provides built-in consent and preference management capabilities to honor user privacy choices. It is GDPR and CCPA compliant, offering features to track user consent status and immediately suppress or delete data when a user revokes permission. For example, if a customer opts out or withdraws consent, Blueshift can automatically exclude that profile from activations (campaigns, syndications) going forward. The platform has a notification preference center (standard out-of-the-box for unsubscribe/opt-down management), and it logs "Notification Preference" changes in the audit trail, indicating those changes take effect across all marketing outputs. Blueshift's documentation specifically notes the "ability to ensure the deletion and automatic suppression of customer data per customer's request" to comply with privacy laws. Therefore, consent changes are instantly respected in all downstream engagement.	https://help.blueshift.com/hc/en-us/articles/ https://blueshift.com/plans-and-pricing/ https://help.blueshift.com/hc/en-us/articles/

Oct 2025 CDP Pulse Report - Blueshift

Yes/No	Details	Source(s)
13. Does your solution support regional data residency options (e.g., EU-only, US-only) with encryption at rest (AES-256) and in transit (TLS 1.2+)?	Yes Blueshift offers regional data residency and strong encryption. The company maintains separate data centers in the USA and EU, and customers can choose to have all their data processed in the EU instance if needed (ensuring data does not leave that region). Moreover, Blueshift adheres to enterprise security standards: data is encrypted at rest and in transit. Blueshift confirms that all customer data is encrypted using AES-256 for storage and TLS 1.2+ for data in transit. It is also SOC 2 Type II certified and GDPR compliant. In summary, you can opt for EU-only or US-only hosting, and in all cases Blueshift protects data with encryption both at-rest and in-transit by default.	https://help.blueshift.com/hc/en-us/articles/4404624313235-Blueshift-global-regions
14. Does your solution provide reverse-ETL / audience-syndication connectors that push profiles and events to major ad, analytics, and messaging destinations natively?	Yes Blueshift provides extensive native audience syndication and reverse-ETL capabilities directly to major advertising, analytics, and messaging platforms. Additionally, it offers integrations with external reverse ETL providers like Census and RudderStack, allowing further flexibility and expanded connectivity options when desired.	https://help.blueshift.com/hc/en-us/articles/4405437312915-Privacy-compliance
15. Does your solution expose a unified customer-profile API/stream that any downstream system can query or subscribe to in real time?	Yes Blueshift exposes APIs for unified customer profiles and events so external systems can retrieve or subscribe to data in real time. For example, Blueshift's REST API includes endpoints to search and fetch customer profiles (by ID or attributes) on-demand. This allows any downstream app (like a customer service system or custom app) to query the CDP for the latest profile data. Additionally, Blueshift can stream data out via its integrations – e.g. it supports real-time event forwarding to webhook-based "Custom Apps" or partners, and incremental audience syncs. Once data is ingested, it's immediately available for use in journeys or to be sent outward. Blueshift's documentation emphasizes that profiles are updated in real time and can be used to personalize external touchpoints. It even offers a "Profile summary API for contact center personalization" as part of its enterprise features. In summary, developers can access a unified profile through Blueshift's API and set up real-time data feeds to other systems, fulfilling this requirement.	https://developer.blueshift.com/reference/post_live https://blueshift.com/plans-and-pricing/
16. Does your solution publish transparent, usage-based pricing (events, MTU, storage) in writing, independent of professional-services spend?	Yes Blueshift provides transparent, usage-based pricing information on its website. The pricing page outlines tiers and included usage limits for profiles, events, and messages. For example, the Growth plan (CDP) "starts at \$750/month" (annual billing) and includes up to 100,000 profiles and 2 million events per month in the base package. It also shows how higher custom volumes can be accommodated in the Enterprise plan. These published plans indicate costs scale with metrics like number of tracked events, profiles, and messaging volume – independent of any services. Blueshift even offers a free trial tier (100k events, 10k profiles free), demonstrating the usage-based model. All pricing details (subscription fees for software usage) are provided in writing on their site, separate from any professional services, which are not required to utilize the platform.	https://blueshift.com/plans-and-pricing

Vendor Deep Dive



Marigold

Primary Archetype: Engagement Hub

Strengths: Retail Specialist

Capability Lens: In-Session Activation

VOX Score: 62.38 / 100

Preference Coverage: 10 / 16 questions “Yes”

📞 Talk to Marigold: <https://meetmarigold.com>

📞 [Ask VOX for Introduction](#)

Oct 2025 CDP Pulse Report - Marigold

Yes/No	Details	Source(s)	
1. Does your solution provide data-observability dashboards that refresh in ≤ 60 seconds (p95) and show source-to-destination lineage for every first-party data stream?	No	Marigold's platform focuses on real-time campaign performance analytics rather than dedicated data observability/lineage dashboards. It offers "real-time insights" for immediate visibility into campaign metrics (e.g. opens, clicks) with fast-refreshing dashboards, but it does not include a specialized interface for end-to-end data pipeline lineage across sources and destinations. There is no evidence in Marigold's documentation of a feature that visualizes source-to-destination data lineage or a guarantee of sub-minute data observability updates for all data streams.	https://meetmarigold.com/platform
2. Does your solution ship native inbound connectors or APIs for the top-10 commerce, marketing-automation, and service platforms without custom code or PS?	Yes	Marigold offers hundreds of pre-built connectors and robust APIs to integrate with popular commerce, marketing, and service platforms out-of-the-box. These native connectors cover top systems (e.g. Shopify and Adobe Commerce for ecommerce, Salesforce and Microsoft Dynamics for CRM/service, popular marketing automation tools, data warehouses, etc.) and enable bi-directional data sync without custom coding or professional services. This means Marigold can natively ingest data from (and send data to) the most widely-used platforms in commerce, marketing automation, and customer service categories, accelerating integration timelines.	https://meetmarigold.com/platform https://meetmarigold.com/industry/retail
3. Does your solution support server-side and client-side collection for web, mobile-app, email, and in-store (POS/scanner) events in a single SDK?	No	Marigold supports both client-side and server-side data collection across channels, though not via a single unified SDK. It provides client-side SDKs (e.g. a JavaScript tag for web tracking and a Mobile SDK for iOS/Android to capture in-app and push notification events) and server-side integration options (open APIs and connectors for uploading events from backend systems or POS data). Together, these allow collection of web behaviors, mobile app interactions, email engagement (tracked natively by the platform), and offline/POS events (via API or scheduled imports). However, these capabilities are delivered through multiple specialized SDKs/APIs rather than one all-in-one library.	
4. Does your solution offer deterministic identity-resolution that stitches online cookies, MAIDs, loyalty IDs and POS receipts into one customer profile?	Yes	Marigold's platform includes a Unified Customer Profiles engine with deterministic identity resolution. It can link multiple identifiers — for example, browser cookies, mobile device IDs (MAIDs), loyalty program IDs, in-store/POS transaction identifiers (receipts or member IDs) — into a single profile. The identity resolution uses a graph-based mapping of identifiers and a rules engine to merge aliases under a persistent customer ID. This ensures that even as customers engage across online and offline touchpoints, Marigold can stitch those interactions together into one unified profile (with full auditability and privacy compliance built-in).	https://meetmarigold.com/platform https://meetmarigold.com/platform
5. Does your solution sync online and in-store inventory events in < 5 minutes (p95) to enable BOPIS and endless-aisle use cases?	Yes	Marigold can handle near-real-time syncing of online and offline data, suitable for BOPIS/endless aisle. The platform's connector framework allows frequent or continuous updates from inventory systems. For example, native integrations with e-commerce platforms (like Shopify or Adobe Commerce) and POS/CRM systems can import events (e.g. product stock changes or store sales) on a minute-level schedule. Marigold's Data Sync can be scheduled in short intervals (even single-digit minutes), and the system is designed for real-time activation of data. While Marigold's marketing content doesn't explicitly quote a 5-minute SLA for inventory, its real-time data ingestion and event-trigger capabilities suggest inventory updates would be reflected within minutes. This enables use cases like Buy Online Pickup In-Store (BOPIS) and endless aisle by keeping product availability in sync across channels with low latency.	https://meetmarigold.com/industry/retail https://getstarted.meetmarigold.com/engage/

Oct 2025 CDP Pulse Report - Marigold

	Yes/No	Details	Source(s)
6. Does your solution expose decision or personalization APIs that return in <1 second (p95) so actions fire while the session is active?	Yes	Marigold provides real-time decisioning and personalization capabilities through its APIs. The platform can respond to inbound triggers and API calls within sub-second latency, enabling in-session personalization. For instance, Marigold's "Real-time adaptability" feature allows content or messages to be updated on-the-fly "in milliseconds" as user behavior changes. Additionally, Marigold's open architecture (including RESTful APIs) supports querying customer profiles or getting personalized recommendations dynamically during an active session. This means that if a downstream system (e.g. a website or mobile app) calls Marigold for a decision (such as next-best offer or experience) while a user is live, the API will return fast enough (typically under 1 second) to execute the action before the user moves on.	https://meetmarigold.com/solutions/activate
7. Does your solution include embedded predictive-merchandising models (e.g., next-best product) deployable by business users without data-science help?	Yes	Marigold comes with built-in predictive analytics and merchandising models that marketers can use without coding. The platform's AI/ML features include product and offer recommendations based on customer behavior. For example, Marigold's predictive models can determine optimal content or offers per user: one out-of-the-box capability is "Predictive offer optimization" where machine learning predicts the best discount or incentive for each customer. Similarly, Marigold (via its Selligent and Sailthru heritage) offers automated product recommendations (like next-best product or content) using algorithms that analyze browsing and purchase history. These tools are accessible in the campaign interface (often under "Smart Content" or personalization blocks) so that a marketer can deploy next-best-product recommendations in emails or on the website without needing a data scientist – simply by configuring the desired algorithm (e.g. collaborative filtering, trending products) in the platform.	https://meetmarigold.com/platform https://getstarted.meetmarigold.com/engage/ https://getstarted.meetmarigold.com/engage/
8. Does your solution allow teams to upload and deploy their own ML/AI models (PMML/ONNX/Python) to score profiles or events in real time?	No	Marigold does not support customer-uploaded custom ML models in its standard offering. Users cannot directly import or deploy arbitrary external ML models (e.g. PMML, ONNX, or custom Python code) into the Marigold platform for real-time scoring. Instead, Marigold focuses on its built-in AI capabilities (such as its native predictive scoring, recommendations, and AI-driven segmentation). While these built-in models cover common use cases, the platform does not advertise any feature for uploading custom-trained machine learning models to run within the system's real-time workflow. Organizations requiring custom model scoring would typically have to score data outside Marigold and then import the results via Marigold's APIs or connectors.	
9. Does your solution provide an SQL workspace or direct lake-house access for segmentation, attribution, and ad-hoc analytics?	Yes	Marigold enables advanced analytics through both an embedded querying capability and data lake integrations. For marketers, the platform offers a flexible, relational data model for segmentation (including an Audience builder and even SQL-like expressions for advanced users in some modules). While Marigold's UI is geared toward no-code segmentation, power users can access data via Marigold's APIs or export connectors. In practice, Marigold provides direct data access options: it has pre-built connectors to data warehouses and lakes (e.g. Snowflake, BigQuery, Amazon Redshift), allowing clients to sync all their customer data to a lake-house for ad-hoc SQL analysis or to connect BI tools. The platform also supports an SQL-based segment definition ("SQL containers") internally for complex segmentation logic in Marigold Engage (for advanced marketers comfortable with SQL). These options mean teams can perform custom queries for segmentation, attribution modeling, and other analytics either within the platform or by leveraging their own data lake connected to Marigold.	https://meetmarigold.com/industry/retail

Oct 2025 CDP Pulse Report - Marigold

	Yes/No	Details	Source(s)
10. Does your solution run continuous data-quality and anomaly detection with automatic quarantine or alerting before bad data feeds downstream systems?	No	<p>Marigold does not have a specific built-in “data observability” module for automated anomaly detection or quarantine of bad data prior to downstream use. The platform ensures data integrity through governance and monitoring practices (e.g. it logs and audits all key data changes, and it offers integrations for email validation to improve data quality), but it relies on configuration and external processes rather than an automatic quarantine system. There is no mention in Marigold’s documentation of real-time anomaly detection on incoming data streams or a feature that automatically halts “bad” data. Instead, data quality is maintained via defined validation rules and human oversight. Alerts can be set up for integration failures (e.g. if a scheduled data sync fails), but a proactive anomaly detection (e.g. detecting out-of-range values and stopping them) is not described as part of the core product.</p>	
11. Does your solution enforce role-based access controls and immutable audit logs at the field level across all regions?	Yes	<p>Marigold implements robust security and compliance controls including role-based access control (RBAC) and audit logging. Administrators can define user roles and access groups to restrict who can view or modify specific data and features (for example, advanced reports are by default only accessible to admins unless permissions are granted to other roles). Every important action in the platform is centrally logged and auditable – Marigold notes that “all key actions on the application are centrally logged, audited and monitored,” and even internal support access to customer accounts is recorded for audit purposes. While the audit logs are not exposed for end-user editing (ensuring immutability), they provide a compliance trail of changes at both configuration and data level. Field-level permission control (masking or restricting certain data fields) can be achieved via the platform’s permission settings and data policies, which is important for multi-region compliance. Marigold holds SOC 2 certification, reflecting these stringent access and audit controls across its global infrastructure.</p>	https://getstarted.meetmarigold.com/engageplus/ https://www.campaignmonitor.com/trust/security/
12. Does your solution offer built-in consent & preference management that suppresses activation instantly when permissions change (GDPR/CCPA ready)?	Yes	<p>Marigold includes built-in consent and preference management tools to honor user privacy choices. The platform has native features for managing subscription statuses, consent flags, and communication preferences which are directly integrated into campaign audience selection. If a customer updates their consent (for example, opts-out or changes preferences), Marigold immediately reflects that change and will suppress that individual from any activation or messaging that they no longer consent to. Marigold (especially through its Selligent module) is designed to be GDPR/CCPA compliant: it allows dynamic consent management and has preference centers that brands can use. In fact, Marigold highlights providing “built-in consent and preference tools aligned with European legislation,” enabling instant suppression of communications when permissions are withdrawn. This ensures no messages are sent in violation of a user’s latest consent status.</p>	https://www.selligent.com/resources/blog/
13. Does your solution support regional data residency options (e.g., EU-only, US-only) with encryption at rest (AES-256) and in transit (TLS 1.2+)?	Yes	<p>Marigold supports regional data residency and employs strong encryption. The company operates data centers in multiple regions (for instance, data is stored in the U.S. by default, with additional data processing infrastructure in regions like Germany (EU) and Australia). Enterprise clients can work with Marigold to ensure data is kept in-region to meet EU-only or other residency requirements (Marigold’s Selligent unit has a large European presence for EU data handling). On security, Marigold meets modern encryption standards: all data in transit is protected via TLS 1.2+ secure protocols, and data at rest in its databases is encrypted using AES-256 encryption. These measures apply across all regions where Marigold stores data. Marigold’s compliance certifications (like ISO 27001, SOC 2, etc.) further underscore its commitment to regional data security and privacy.</p>	https://www.campaignmonitor.com/trust/security/ https://www.campaignmonitor.com/trust/security/

Oct 2025 CDP Pulse Report - Marigold

Yes/No	Details	Source(s)	
14. Does your solution provide reverse-ETL / audience-syndication connectors that push profiles and events to major ad, analytics, and messaging destinations natively?	No	Marigold's platform focuses on ingesting and activating data within its own ecosystem, and it does not natively offer a broad "reverse-ETL" module for arbitrarily syncing audiences out to external ad platforms or analytics tools in the way specialized CDPs do. While Marigold has bi-directional connectors for certain systems (e.g. syncing data to CRM or exporting to data warehouses), it does not list out-of-the-box connectors to push customer audiences directly into major advertising networks (such as Google Ads Customer Match or Facebook Custom Audiences) or into third-party analytics SaaS. Any such syndication would typically require using Marigold's APIs or a third-party integration (for example, customers might export segments to CSV/warehouse and then use their own tools or iPaaS to upload to ad platforms). In summary, Marigold can export data (profiles or event lists) via its connectors and API, but it does not have a native one-click audience syndication feature aimed at marketing destinations like Google, Facebook, etc.	https://meetmarigold.com/industry/retail
15. Does your solution expose a unified customer-profile API/stream that any downstream system can query or subscribe to in real time?	Yes	Marigold provides a unified customer profile accessible via API, and supports real-time data streaming out. The platform's architecture includes RESTful APIs that allow external systems to query profile data on demand (for example, retrieving a customer's attributes or segment memberships). Marigold highlights having "REST-based APIs" and profiles that are "updated in real-time," which indicates that an external system can pull the latest profile info programmatically at any time. Additionally, Marigold can stream events out in real-time: for instance, it has a connector to publish events to Google Pub/Sub and can act as a SegmentDestination (receiving unified data from Segment, or conversely sending data out). These capabilities mean you can subscribe a downstream system to changes (either by polling the profile API or by receiving event streams/webhooks). In practice, if a customer takes an action, Marigold can immediately update their unified profile and an external tool can fetch or get notified of that update in real-time.	https://meetmarigold.com/solutions/messaging/sailthru https://getstarted.meetmarigold.com/platform/
16. Does your solution publish transparent, usage-based pricing (events, MTU, storage) in writing, independent of professional-services spend?	No	Marigold does not publicly disclose transparent, usage-based pricing on its website for its enterprise offerings. Pricing for Marigold's solutions (such as Cheetah Digital, Selligent, etc.) is typically handled via custom quotes and is not broken out in publicly available rate cards. There is no self-service pricing calculator listing costs per event, MAU/MTU, or data volume. Moreover, any pricing information is not independent of service engagements – often the licensing and any needed onboarding services are discussed together in sales engagements. In contrast to some newer SaaS CDPs, Marigold (being an enterprise marketing platform) does not have its usage-based pricing published in writing. Even on software comparison sites, Marigold's pricing is listed as "No information available." Prospective customers must contact Marigold for a tailored pricing proposal.	https://sourceforge.net/software/compare/

Vendor Deep Dive



Optimove

Primary Archetype: Engagement Hub

Strengths: Retail Specialist

Capability Lens: In-Session Activation

VOX Score: 74.77 / 100

Preference Coverage: 12 / 16 questions “Yes”

📞 Talk to Optimove: www.optimove.com

📞 [Ask VOX for Introduction](#)

Oct 2025 CDP Pulse Report - Optimove

	Yes/No	Details	Source(s)
1. Does your solution provide data-observability dashboards that refresh in ≤ 60 seconds (p95) and show source-to-destination lineage for every first-party data stream?	No	Optimove's documentation does not indicate a dedicated data observability or lineage dashboard with 60-second refresh. The platform offers customizable dashboards and a BI Studio for monitoring marketing KPIs, but there is no mention of real-time (≤60s) data pipeline freshness or end-to-end data lineage visualization for each data stream.	
2. Does your solution ship native inbound connectors or APIs for the top-10 commerce, marketing-automation, and service platforms without custom code or PS?	Yes	Optimove provides many out-of-the-box integrations and APIs for popular marketing, commerce, and service systems. Its built-in connectors cover major email service providers, marketing automation tools, ad networks, etc. (e.g. Adobe Campaign, Marketo, Salesforce Marketing Cloud, Facebook Ads, Google Ads) without requiring custom code. Additionally, a rich Optimove API allows integrating any other in-house or third-party systems as needed.	https://www.optimove.com/faq/how-does-optimove-integrate-with-my-marketing-and-campaign-management-systems https://www.optimove.com/faq/how-does-optimove-integrate-with-my-marketing-and-campaign-management-systems
3. Does your solution support server-side and client-side collection for web, mobile-app, email, and in-store (POS/scanner) events in a single SDK?	Yes	Optimove supports both client-side and server-side data collection across channels. It offers a Web SDK for tracking website user behavior and events, native Mobile SDKs for apps, and also allows server-side event ingestion via APIs or batch uploads (for sources like POS or offline data). All these feeds are unified into the Single Customer View, enabling collection of web, mobile, email engagement, and offline purchase events into one platform.	https://developer.optimove.com/docs/web-sdk-integration https://www.optimove.com/faq/how-does-optimove-integrate-with-my-marketing-and-campaign-management-systems
4. Does your solution offer deterministic identity-resolution that stitches online cookies, MAIDs, loyalty IDs and POS receipts into one customer profile?	Yes	Optimove's Customer Data Platform builds a unified 360° Single Customer View by stitching together identifiers from multiple sources. The platform integrates first-party raw data, on-site/in-app activity, third-party data, and offline/campaign data into one customer profile. This implies deterministic matching of identifiers (e.g. linking online identities like cookies or mobile IDs with offline IDs like loyalty or POS transaction IDs) to consolidate all data on each customer.	https://www.optimove.com/product/optimove-orchestrate
5. Does your solution sync online and in-store inventory events in < 5 minutes (p95) to enable BOPIS and endless-aisle use cases?	No	There is no clear evidence that Optimove handles inventory data synchronization within 5 minutes for BOPIS/endless aisle. Optimove's real-time capabilities center on customer events and marketing triggers (leveraging an event-streaming infrastructure with Apache Kafka) rather than inventory management. The documentation does not mention any specific feature for rapid (<5min) inventory stock updates across online and in-store systems.	
6. Does your solution expose decision or personalization APIs that return in <1 second (p95) so actions fire while the session is active?	Yes	Optimove exposes real-time decisioning and personalization through its APIs/SDK. For example, it provides a Recommendation API for personalized content (product recommendations) that can be called on-the-fly. The platform's architecture (using Confluent/Kafka event streaming) is built for low-latency responses, enabling in-session triggers like web personalization, product recommendations, or next-best-action decisions to execute within sub-second response times while the user is still active.	https://www.optimove.com/industries/retail https://www.optimove.com/product/optimove-orchestrate
7. Does your solution include embedded predictive-merchandising models (e.g., next-best product) deployable by business users without data-science help?	Yes	Optimove includes built-in AI predictive models for merchandising and next-best offers that marketers can use without coding. Its platform comes with over 20 AI-driven recommendation models for personalized product suggestions (e.g. "Similar Items," "Trending Near You," "Because You Bought"). Optimove's "Optibot" and AI features also enable next-best-action or next-best-offer recommendations and Self-Optimizing campaigns out-of-the-box, allowing business users to deploy predictive merchandising strategies without data science involvement.	https://www.optimove.com/product/optimenie-ai https://www.optimove.com/blog/

Oct 2025 CDP Pulse Report - Optimove

Yes/No	Details	Source(s)
8. Does your solution allow teams to upload and deploy their own ML/AI models (PMML/ONNX/Python) to score profiles or events in real time?	No	Optimove does not advertise any functionality for users to directly import or deploy custom user-trained ML models (e.g., PMML, ONNX, or custom Python models) for real-time scoring within the platform. Instead, Optimove provides predictive customer models and AI algorithms developed internally by its data science team, tailored individually for each client. While there's no self-service capability for full-model uploads, Optimove does offer custom model development and implementation through its Professional Services team. Additionally, scores generated externally by client-trained models can be uploaded into Optimove and updated in real-time as standard customer data attributes.
9. Does your solution provide an SQL workspace or direct lake-house access for segmentation, attribution, and ad-hoc analytics?	Yes	Optimove enables direct data access for advanced analytics. It leverages Snowflake Data Sharing to give clients a "Data Share" – a repository of unified, cleansed customer and event data that can be queried with SQL. This means organizations can connect to Optimove's Snowflake-based data share and run ad-hoc SQL queries or use BI tools on their customer data (for custom segmentation, attribution analysis, etc.) outside of the Optimove UI, essentially providing a lakehouse-style access to the CDP data.
10. Does your solution run continuous data-quality and anomaly detection with automatic quarantine or alerting before bad data feeds downstream systems?	No	Optimove's publicly available information does not describe an automated data quality monitoring system that quarantines bad data. The platform does assist with data cleansing and de-duplication during data ingestion (using proprietary algorithms per each client's policies). Additionally, its AI assistant (Optibot) can identify anomalies in campaign performance or data trends. However, there is no mention of continuous anomaly detection on incoming data with automatic quarantine or alerts to stop bad data before it propagates to outputs.
11. Does your solution enforce role-based access controls and immutable audit logs at the field level across all regions?	Yes	Optimove follows enterprise-grade security practices, including role-based access control and auditing. The platform is SOC 2 Type II and ISO 27001 certified, indicating strict controls and audit processes. Optimove implements a zero-trust security model and monitors all access; for example, it has a Security Operations Center that receives alerts on any suspicious access/log anomalies. While field-level RBAC granularity isn't explicitly detailed in public docs, these certifications and practices imply that Optimove maintains comprehensive, immutable audit logs and region-specific compliance (e.g., meeting HIPAA, GDPR requirements for data access).
12. Does your solution offer built-in consent & preference management that suppresses activation instantly when permissions change (GDPR/CCPA ready)?	Yes	Optimove provides built-in consent and preference management tools. It includes a Preference Center module for managing customer communication preferences (opt-ins/opt-outs). The platform's APIs allow retrieving and updating these preferences and unsubscribe lists in real time. When a customer's consent status changes (e.g. they withdraw consent or unsubscribe), Optimove will honor that immediately – ensuring that marketing activations are suppressed for that user as required. Optimove's adherence to GDPR is explicitly stated, and its features support compliance by allowing instant suppression based on consent changes.

Oct 2025 CDP Pulse Report - Optimove

Yes/No	Details	Source(s)
13. Does your solution support regional data residency options (e.g., EU-only, US-only) with encryption at rest (AES-256) and in transit (TLS 1.2+)?	Yes Optimove is designed with strong data security and privacy compliance, suggesting support for regional data residency. They comply with GDPR and other global privacy frameworks (EU-US Privacy Shield, etc.), which implies accommodating EU-only data requirements when needed. In terms of security, Optimove meets leading standards (ISO 27001, SOC 2) and employs best practices like encryption of data in transit and at rest. While exact encryption protocols aren't listed on the marketing site, these certifications mean Optimove uses industry-standard encryption (e.g., TLS 1.2+ for data in motion and AES-256 for data at rest) to protect data. They also have a full disaster recovery site, indicating robust data protection across regions.	https://www.optimove.com/blog/7-tips-to-better-align-your-marketing-and-it-teams-when-evaluating-new-technologies
14. Does your solution provide reverse-ETL / audience-syndication connectors that push profiles and events to major ad, analytics, and messaging destinations natively?	Yes Optimove natively supports pushing customer data and audiences to many external destinations. It offers 50+ pre-built execution channel integrations covering major ad platforms, analytics, and messaging tools. For instance, Optimove can directly synchronize segments or campaign audiences to Facebook Custom Audiences, Google Ads/GA, email/SMS providers, etc., through its partner integrations. It also has the ability to export data via API or file for any other destinations. This means profiles and event data can be automatically syndicated to downstream advertising, analytics, or communication systems without custom development.	https://www.optimove.com/blog/7-tips-to-better-align-your-marketing-and-it-teams-when-evaluating-new-technologies https://www.optimove.com/faq/how-does-optimove-integrate-with-my-marketing-and-campaign-management-systems
15. Does your solution expose a unified customer-profile API/stream that any downstream system can query or subscribe to in real time?	Yes Optimove provides a unified customer-profile API that downstream systems can query in real-time, such as the GET /Customers/{ccid}/Attributes endpoint, which returns current customer attribute values on demand. Additionally, Optimove supports real-time event notifications through webhooks and event-listener APIs, allowing external systems to subscribe and receive updates about customer events, campaign eligibility, and profile changes.	https://developer.optimove.com/reference/get_customers-ccid-attributes
16. Does your solution publish transparent, usage-based pricing (events, MTU, storage) in writing, independent of professional-services spend?	Yes Optimove's pricing model is described transparently and is usage-based. According to their FAQ, the core platform pricing is based on the number of monthly active customers (an MTU-like metric) that the client has. They do not charge extra for the number of user seats or for data volume shared. Add-on products are priced by usage as well – for example, Optimove's email sending add-on is priced by the volume of emails sent. This pricing information is documented in writing and is separate from any professional services, which are available but not required (professional services are offered as optional support and are not tied to the licensing cost of the software).	https://www.optimove.com/faq/how-does-optimove-integrate-with-my-marketing-and-campaign-management-systems https://www.optimove.com/faq/how-does-optimove-integrate-with-my-marketing-and-campaign-management-systems

Suite CDPs



Vendor Deep Dive



mParticle

Primary Archetype: Suite CDP

Strengths: Warehouse-native/Composable;
Governance & Residency

Capability Lens: Data Visibility & Quality; Security, Consent & Residency; Customization & Composability

VOX Score: 80.86 / 100

Preference Coverage: 13 / 16 questions “Yes”

📞 Talk to mParticle: www.mparticle.com

📞 [Ask VOX for Introduction](#)

Oct 2025 CDP Pulse Report - mParticle

Yes/No	Details	Source(s)
1. Does your solution provide data-observability dashboards that refresh in ≤ 60 seconds (p95) and show source-to-destination lineage for every first-party data stream?	Yes mParticle offers real-time data observability. The platform's Live Stream provides a real-time view of all data flowing in and out, updating within seconds. Additionally, Observability traces in mParticle show end-to-end data lineage – i.e. exactly where data came from and where it went across sources and destinations. These tools let users quickly spot issues (within ~seconds) and visualize source-to-destination flows for each data stream.	https://docs.mparticle.com/guides/platform-guide/live-stream/ https://docs.mparticle.com/guides/platform-guide/observability/overview/
2. Does your solution ship native inbound connectors or APIs for the top-10 commerce, marketing-automation, and service platforms without custom code or PS?	Yes mParticle provides hundreds of pre-built integrations out-of-the-box. It natively supports feeds (inbound connectors) for major commerce platforms (e.g. Shopify), marketing automation tools, and service/CRM systems (e.g. Salesforce, Zendesk) – all configurable without custom coding. In total, mParticle connects to “hundreds of vendors” via SDKs/APIs, simplifying data collection from virtually all top systems without needing professional services.	https://docs.mparticle.com/guides/platform-guide/observability/guide/ https://support.mparticle.com/hc/en-us/articles/11368255468941-mParticle-Integration
3. Does your solution support server-side and client-side collection for web, mobile-app, email, and in-store (POS/scanner) events in a single SDK?	Yes mParticle supports multi-channel data collection across client and server. It offers Web, iOS, Android and other client SDKs, plus server-side APIs/SDKs (Node, Python, etc.) for backend or POS data. All these inputs feed into one unified platform – mParticle “collects data from all of your platforms” (web, mobile, email, in-store) and stitches it together centrally. In practice, developers use the appropriate SDK for each environment, but the data converges into a single pipeline.	https://docs.mparticle.com/guides/platform-guide/tiered-events/ https://docs.mparticle.com/guides/platform-guide/tiered-events/
4. Does your solution offer deterministic identity-resolution that stitches online cookies, MAIDs, loyalty IDs and POS receipts into one customer profile?	Yes mParticle performs deterministic identity resolution to unify user identities. It uses a framework called IDSync to rules-match and merge identifiers (e.g. device cookie IDs, mobile Ad IDs (MAIDs), loyalty or CRM IDs, offline transaction IDs) into one mParticle ID (MPID) profile. The platform “resolves customer identities and stitch[es] together customer profiles” across touchpoints. This deterministic approach ensures online and offline identifiers for the same person are linked to a single unified customer profile.	https://docs.mparticle.com/guides/platform-guide/tiered-events/ https://docs.mparticle.com/guides/idsync/identify-users/
5. Does your solution sync online and in-store inventory events in < 5 minutes (p95) to enable BOPIS and endless-aisle use cases?	Yes mParticle's event pipeline operates in real time, easily meeting a sub-5-minute latency for inventory updates. The system streams events as they're received – providing a “real time view of all data” and forwarding to outputs almost immediately. In practice, inventory changes (whether from online e-commerce or in-store POS) are ingested and forwarded in seconds, enabling near-instant data sync for BOPIS/endless aisle scenarios. (mParticle's architecture is built for real-time data streaming.)	https://docs.mparticle.com/guides/platform-guide/live-stream/ https://www.cmswire.com/digital-experience/4-cdp-case-studies-openprise-tealium-segment-mparticle/
6. Does your solution expose decision or personalization APIs that return in < 1 second (p95) so actions fire while the session is active?	Yes mParticle provides real-time personalization APIs, notably the Profile API, to support in-session decisions. The Profile API delivers a user's unified profile (attributes, segments, etc.) in JSON, enabling personalized experiences on the fly. This API is designed for sub-second responses. Additionally, mParticle's Predictive Audience features (e.g. Next Best Action) compute recommendations that can be accessed in real time. Together, these allow external systems (web or mobile apps) to query mParticle for personalization data and get a response fast enough (<1s) to act during the active session.	https://docs.mparticle.com/guides/customer-360/profiles/overview/ https://docs.mparticle.com/guides/customer-360/profiles/overview/
7. Does your solution include embedded predictive-merchandising models (e.g., next-best product) deployable by business users without data-science help?	Yes mParticle offers no-code predictive modeling for marketers via its Cortex AI (Vidora) integration. After acquiring Vidora, mParticle enables teams to “easily create and deploy models” like next-best-action/offers without coding. Business users can configure predictions (churn risk, next product to buy, etc.) through the UI as Predictive Attributes or Next Best Action strategies. The system was built to be self-service for business users, so data science resources are not required for using these embedded machine learning models.	https://www.permira.com/news-and-insights/announcements/mparticle-acquires-ai-startup-vidora-to-help-teams-augment-customer-profiles-and-improve-their-marketing-stacks https://www.permira.com/news-and-insights/announcements/mparticle-acquires-ai-startup-vidora-to-help-teams-augment-customer-profiles-and-improve-their-marketing-stacks

Oct 2025 CDP Pulse Report - mParticle

	Yes/No	Details	Source(s)
8. Does your solution allow teams to upload and deploy their own ML/AI models (PMML/ONNX/Python) to score profiles or events in real time?	No	mParticle does not support uploading custom ML models into its platform. It provides its own AI modeling (as above) but does not have a feature to ingest a user's PMML/ONNX or custom Python model for native execution. In practice, clients who build custom models apply them outside mParticle and then import the scored results. (As an external analysis notes, if you have your own data science models in a warehouse, bringing those outputs into mParticle requires custom integration and is not handled natively.) There is no built-in interface to deploy arbitrary user-trained models for real-time scoring within mParticle.	https://hightouch.com/blog/what-is-mparticle-cdp https://hightouch.com/blog/what-is-mparticle-cdp
9. Does your solution provide an SQL workspace or direct lake-house access for segmentation, attribution, and ad-hoc analytics?	No	mParticle enables flexible analytics, but not via a built-in SQL editor. Instead, it offers direct data warehouse integrations and an analytics UI. Users can export or sync data to their own Snowflake/BigQuery (lakehouse) and run SQL there, or use mParticle's Indicative-powered interface for segmentation and funnel analysis. mParticle recently introduced Composable Audiences, which let you define audiences by running SQL on your warehouse data (mParticle will execute the queries on your behalf). So while you don't write SQL inside the mParticle app, the platform provides pathways to query and analyze data (via your warehouse or built-in no-code tools for segmentation and attribution).	https://docs.mparticle.com/guides/composable-audiences/ https://www.cdpinstitute.org/wp-content/uploads/2022/01/AU026P-mParticle-RealCDP-2022-Package.pdf
10. Does your solution run continuous data-quality and anomaly detection with automatic quarantine or alerting before bad data feeds downstream systems?	Yes	mParticle has robust data quality controls and alerting. It includes a Data Planning module that enforces schema and can automatically block "bad" events (e.g. schema violations) from being forwarded. The platform's Observability suite also detects anomalies like unexpected drops or surges in event volume, and System Alerts flag any errors in data flows. Invalid events can be routed to a quarantine (e.g. an S3 "Blocked Events" bucket) instead of to downstream tools. Combined, these features ensure anomalous or bad data is caught and alerted on before it pollutes any outputs.	https://docs.mparticle.com/guides/platform-guide/observability/overview/ https://www.cdpinstitute.org/wp-content/uploads/2022/01/AU026P-mParticle-RealCDP-2022-Package.pdf
11. Does your solution enforce role-based access controls and immutable audit logs at the field level across all regions?	Yes	mParticle enforces strict role-based access control (RBAC) and maintains detailed logs of system access and changes. The platform has granular user roles/permissions (including custom roles) and follows a least privilege approach – even mParticle's internal staff have no default access to customer data. All production access is gated by roles, and activity is audited regularly. While field-level restrictions are not commonly used in the UI, sensitive data can be masked or restricted via privacy rules. mParticle's security documentation confirms periodic access audits and immutable logging of critical actions for compliance.	https://www.indicative.com/terms-of-service-08-16-22/ https://www.indicative.com/terms-of-service-08-16-22/
12. Does your solution offer built-in consent & preference management that suppresses activation instantly when permissions change (GDPR/CCPA ready)?	Yes	mParticle provides native consent and preference management as part of its Data Privacy Controls. It allows you to define consent purposes (e.g. GDPR marketing consent or CCPA "do not sell") and store each user's consent state. That consent state is then automatically enforced in real time – for example, you can configure mParticle to "only forward user data if GDPR consent for 'marketing' is true." If a user opts out or withdraws consent, mParticle immediately stops sending their data to any disallowed outputs. This built-in framework covers GDPR/CCPA requirements and ensures downstream activations always respect the latest user permissions.	https://docs.mparticle.com/guides/data-privacy-controls/ https://docs.mparticle.com/guides/data-privacy-controls/
13. Does your solution support regional data residency options (e.g., EU-only, US-only) with encryption at rest (AES-256) and in transit (TLS 1.2+)?	Yes	mParticle supports data locality and strong encryption. Customers can choose to host data in localized regional data centers – for example, EU or Australia pods instead of US – to meet data residency needs. All data is secured both in transit and at rest: mParticle uses TLS for all data transport and employs AES-256 encryption for data at rest in its AWS and GCP environments. These measures align with industry standards (TLS 1.2+ and AES-256) and ensure compliance with regional data protection regulations.	https://docs.mparticle.com/developers/guides/data-localization/ https://www.indicative.com/terms-of-service-08-16-22/
14. Does your solution provide reverse-ETL / audience-syndication connectors that push profiles and events to major ad, analytics, and messaging destinations natively?	Yes	mParticle's core functionality is to syndicate customer data to hundreds of downstream systems. It offers 300+ pre-built outbound connectors covering all major advertising platforms, analytics tools, and marketing/messaging services. This includes native integrations to push audience segments or event data to Google Ads, Facebook, Braze, Adobe Analytics, Salesforce Marketing Cloud, and many more. These connectors operate in real time and require no custom development, effectively delivering reverse-ETL capabilities as a built-in feature of the mParticle CDP.	https://www.cdpinstitute.org/wp-content/uploads/2022/01/AU026P-mParticle-RealCDP-2022-Package.pdf https://support.mparticle.com/hc/en-us/articles/11368255468941-mParticle-Integration

Oct 2025 CDP Pulse Report - mParticle

Yes/No	Details	Source(s)	
15. Does your solution expose a unified customer-profile API/stream that any downstream system can query or subscribe to in real time?	Yes	<p>mParticle provides a unified profile API and real-time event streams. The Profile API allows external systems to query a customer's consolidated profile (identifiers, attributes, consents, etc.) on demand in real time. Additionally, mParticle supports real-time forwarding of event streams via webhooks or stream integrations (e.g. AWS Kinesis), which downstream systems can subscribe to. Together, these features mean any authorized system can either call for the latest profile data or subscribe to live event feeds from mParticle, enabling real-time personalization and data syncing based on the unified customer profile.</p>	https://docs.mparticle.com/guides/customer-360/profiles/overview/ https://docs.mparticle.com/guides/platform-guide/live-stream/
16. Does your solution publish transparent, usage-based pricing (events, MTU, storage) in writing, independent of professional-services spend?	No	<p>mParticle's pricing is not publicly transparent. The company does not list standard price tiers or usage rates on its website – interested customers must contact mParticle for a quote. (Industry sources note that mParticle “does not publicly list pricing information” and requires direct inquiry for pricing details.) mParticle has introduced Value-Based Pricing (metered by events, MTUs, etc.), but the specific costs and terms are provided via sales and contracts, not via a public price sheet.</p>	https://panoply.io/compare/grow/mparticle/ https://docs.mparticle.com/guides/platform-guide/platform-settings/key-management/

Vendor Deep Dive



Salesforce Data 360

Primary Archetype: Suite CDP

Strengths: Real-time/Edge; Governance & Residency

Capability Lens: Data Visibility & Quality; Security, Consent & Residency; Customization & Composability

VOX Score: 100 / 100

Preference Coverage: 16/ 16 questions “Yes”

📞 Talk to Salesforce Data 360: <https://www.salesforce.com/data/>

📞 [Ask VOX for Introduction](#)

Oct 2025 CDP Pulse Report - Salesforce Data 360

Yes/No	Details	Source(s)
1. Does your solution provide data-observability dashboards that refresh in \leq 60 seconds (p95) and show source-to-destination lineage for every first-party data stream?	Yes. Salesforce Data 360 includes observability tools like the Data Stream Control Center and Data Cloud Dashboards to monitor data streams in near-real time. Users can view stream status, refresh history, and data quality metrics for first-party data pipelines. Data Cloud's data graph and lineage features allow visualization of relationships between data objects (showing how source data flows into unified profiles and segments). While full end-to-end lineage visualization is limited, the platform does track data source details and offers reports that quickly surface anomalies or errors in data streams (with refresh updates typically within minutes).	https://www.salesforceblogger.com/2024/07/29/gaining-data-cloud-observability-with-the-power-of-reports-and-dashboards/ https://www.salesforceben.com/whats-new-with-data-cloud-sub-second-processing-data-formats-and-connectors/
2. Does your solution ship native inbound connectors or APIs for the top-10 commerce, marketing-automation, and service platforms without custom code or PS?	Yes. Salesforce Data 360 provides 200+ native connectors (including many pre-built integrations) to ingest data from popular commerce, marketing, and service systems with minimal effort. These include connectors for Salesforce's own clouds (Sales, Service, Commerce, Marketing Cloud) and third-party platforms (e.g. Google Ads, Meta, AWS, etc.), as well as a flexible Ingestion API. Many connectors come with starter data mappings, eliminating custom code. Additionally, Salesforce leverages MuleSoft for any less common integrations, ensuring top commerce and marketing platforms can feed data in without professional services.	https://www.salesforceben.com/whats-new-with-data-cloud-sub-second-processing-data-formats-and-connectors/ https://www.salesforce.com/data/connectivity/data-cloud-connectors/
3. Does your solution support server-side and client-side collection for web, mobile-app, email, and in-store (POS/scanner) events in a single SDK?	Yes. Salesforce Data 360 supports both client-side and server-side event collection across channels. It offers Web and Mobile SDKs (Salesforce "Interactions SDK") to capture behavioral events from websites and mobile apps in real time, with built-in consent handling. Email interactions and in-store/POS events can be ingested via native connectors or APIs (e.g. Marketing Cloud for email engagement, and POS data via API or MuleSoft). All these event streams unify in Data Cloud's schema, allowing a single platform/SDK approach to track online and offline customer events together.	https://developer.salesforce.com/blogs/2024/04/using-data-cloud-web-sdk-to-capture-engagement-on-your-website https://www.salesforceben.com/whats-new-with-data-cloud-sub-second-processing-data-formats-and-connectors/
4. Does your solution offer deterministic identity-resolution that stitches online cookies, MAIDs, loyalty IDs and POS receipts into one customer profile?	Yes. Salesforce Data 360 provides deterministic identity resolution to unify identifiers across channels into one profile. It uses configurable matching rules (e.g. email, phone, loyalty ID, device ID) to merge anonymous online identifiers (cookies or device IDs) with known identifiers like loyalty program IDs or in-store purchase records. For example, if a customer logs in or provides an email/loyalty number, Data Cloud immediately merges their web activity with their in-store purchase history into a single profile. This allows cookies, mobile Ad IDs, loyalty accounts, and POS transaction IDs to be linked deterministically, giving a 360° view of the customer.	https://www.salesforce.com/blog/real-time-identity-resolution/ https://www.salesforce.com/blog/real-time-identity-resolution/
5. Does your solution sync online and in-store inventory events in < 5 minutes (p95) to enable BOPIS and endless-aisle use cases?	Yes. Salesforce's platform can sync inventory updates from both online and physical stores in near real-time. Data Cloud's streaming ingestion (sub-minute latency for engagement data) and connectors allow inventory events (stock changes, sales) to flow in roughly within a couple of minutes. This powers BOPIS (Buy Online Pick Up In Store) and Endless Aisle scenarios by ensuring inventory availability is up-to-date across channels. For example, third-party inventory systems can feed Data Cloud (or Salesforce Commerce) and have updates reflected almost instantly for store associates and online shoppers. Salesforce's Retail solutions explicitly emphasize real-time inventory data unification for these use cases.	https://www.cmswire.com/customer-experience/salesforce-introduces-agentforce-for-retail-automation-meets-personalization/ https://www.salesforceben.com/whats-new-with-data-cloud-sub-second-processing-data-formats-and-connectors/
6. Does your solution expose decision or personalization APIs that return in < 1 second (p95) so actions fire while the session is active?	Yes. Salesforce offers real-time decisioning and personalization APIs integrated with Data Cloud. Specifically, Salesforce Personalization (Interaction Studio) provides a Decisioning API that can return personalized content or recommendations within sub-second latency. Salesforce's new "Sub-sec E2E Real-Time" capability in Data Cloud processes inbound events and pairs the unified profile with Einstein personalization models to deliver in-session decisions (e.g. next-best product or offer) well under 1 second. This means websites or mobile apps can call an API and get immediate personalized responses while the user is still engaged.	https://www.salesforce.com/blog/sub-second-real-time/ https://www.salesforce.com/blog/sub-second-real-time/
7. Does your solution include embedded predictive-merchandising models (e.g., next-best product) deployable by business users without data-science help?	Yes. Salesforce provides Einstein AI features that let business users deploy predictive models (like next-best product or offer) with clicks, not code. For example, Einstein Recommendation Builder enables admins/marketers to configure product recommendations or next-best offers using a wizard interface. These models (e.g. upsell suggestions, propensity scores) are embedded in the platform and can be activated by non-technical users through the UI. Additionally, Einstein Next Best Action allows combining these AI predictions with business rules to show recommended products or content to customers – all without requiring a data scientist to build the model from scratch.	https://admin.salesforce.com/blog/2021/get-started-with-einstein-recommendation-builder https://admin.salesforce.com/blog/2021/get-started-with-einstein-recommendation-builder

Oct 2025 CDP Pulse Report - Salesforce Data 360

Yes/No	Details	Source(s)
8. Does your solution allow teams to upload and deploy their own ML/AI models (PMML/ONNX/Python) to score profiles or events in real time?	Yes. With Einstein Studio in Salesforce Data 360, organizations can bring their own models and use them on Data Cloud data in real time. It supports connecting external ML models (e.g. from AWS SageMaker, Databricks, or custom Python/ONNX models) and deploying them to score profiles/events. The imported models can be invoked within Data Cloud for scoring streaming data or profiles (e.g. propensity scores, anomaly detection). For instance, a team can register a pre-trained “propensity to buy” model and Data Cloud will fetch real-time predictions from that model for each customer, which can then trigger personalized actions. This allows custom AI models to run within Salesforce workflows with low latency.	https://developer.salesforce.com/blogs/2023/08/bring-your-own-ai-models-to-salesforce-with-einstein-studio https://developer.salesforce.com/blogs/2023/08/bring-your-own-ai-models-to-salesforce-with-einstein-studio
9. Does your solution provide an SQL workspace or direct lake-house access for segmentation, attribution, and ad-hoc analytics?	Yes. Salesforce Data 360 supports SQL querying and direct data access for advanced analytics. It includes a Query Editor/Calculated Insights feature where users can write ANSI SQL to segment data or create custom metrics. Data Cloud's lakehouse architecture also enables “zero-copy” data shares with external data warehouses (e.g. Snowflake, Amazon Redshift), meaning analysts can access unified profile data in place without moving it. Users can leverage these SQL capabilities for complex segmentation, attribution analysis, or join Data Cloud data with other sources. Additionally, Data Cloud integrates with Tableau and CRM Analytics for BI reporting, further enabling ad-hoc analysis on the unified dataset.	https://www.salesforceben.com/6-ways-to-extract-data-from-salesforce-data-cloud/ https://www.salesforceben.com/what-is-salesforce-data-cloud-vs-salesforce-cdp/
10. Does your solution run continuous data-quality and anomaly detection with automatic quarantine or alerting before bad data feeds downstream systems?	Yes. Salesforce Data 360 includes data quality monitoring and anomaly detection tools to catch issues early. The Data Stream Control Center and Data Cloud Dashboards continually monitor data ingestion and can flag errors or anomalies in incoming data. Salesforce's platform employs anomaly-detection algorithms to identify unusual data patterns in real time (e.g. spikes, missing values). When bad or unexpected data is detected (such as a data stream error), the system can alert users and pause that data from propagating to segments/activations until resolved. While full automatic “quarantine” of bad data may require configuration, Salesforce provides built-in reports and error logs so teams are alerted to data quality issues and can prevent corrupt data from reaching downstream processes.	https://www.salesforceblogger.com/2024/07/29/gaining-data-cloud-observability-with-the-power-of-reports-and-dashboards/ https://www.techcloudsolution.com/blogs/salesforce-data-cloud/
11. Does your solution enforce role-based access controls and immutable audit logs at the field level across all regions?	Yes. Salesforce enforces robust role-based access control (RBAC) and field-level security for Data Cloud data. Admins can define object-level and field-level access policies in Data Cloud, controlling which roles or users can view or edit specific fields. These fine-grained permissions are applied across regions and data spaces to meet compliance needs. In addition, Salesforce maintains immutable audit logs (via tools like Field Audit Trail and Setup Audit Trail) that record changes to data and configuration. Field Audit Trail can retain a history of field value changes and user actions, ensuring an auditable record of who did what with data (and these logs cannot be tampered with). Together, these features provide enterprise-grade security and traceability down to the field level globally.	https://bluecanvas.io/blog/salesforce-role-based-permissions-enhance-security-efficiency-and-collaboration https://help.salesforce.com/s/articleView?id=xcloud.field_audit_trail.htm&language=en_US&type=5
12. Does your solution offer built-in consent & preference management that suppresses activation instantly when permissions change (GDPR/CCPA ready)?	Yes. Salesforce has built-in consent and preference management capabilities that honor customer permissions in real time. Data Cloud integrates Salesforce's consent objects (like Contact Point Consent) to track individuals' GDPR/CCPA consent levels (opt-ins/opt-outs). With the new Contact Point Filtering feature, Data Cloud will automatically filter out or suppress any profile from activations if the required consent isn't present. In practice, this means if a user withdraws consent or changes preferences, that change is immediately respected – the customer's data will not be activated to marketing or advertising channels going forward. This native consent-aware filtering ensures compliance (e.g. no sending to a user who opted out), without requiring custom work, and updates take effect instantly across all downstream engagements.	https://www.salesforceblogger.com/2024/08/26/contact-point-filtering-the-missing-piece-in-your-consent-aware-data-strategy/ https://www.salesforceblogger.com/2024/08/26/contact-point-filtering-the-missing-piece-in-your-consent-aware-data-strategy/
13. Does your solution support regional data residency options (e.g., EU-only, US-only) with encryption at rest (AES-256) and in transit (TLS 1.2+)?	Yes. Salesforce supports regional data residency and strong encryption. Hyperforce, Salesforce's infrastructure architecture, offers data deployment in specified regions (including an EU Operating Zone) to keep customer data in-region for compliance. Data Cloud data is encrypted at rest using AES-256 encryption (Salesforce Platform Encryption), often with customer-managed keys. All data in transit is protected via TLS 1.2+ secure protocols. These measures ensure that an organization can choose where its data is stored (EU-only, US-only, etc.) and that the data is safeguarded by industry-standard encryption both at rest and in motion.	https://www.salesforce.com/platform/data-residence-eu-oz/ https://trailhead.salesforce.com/content/learn/modules/platform-encryption-for-data-cloud/get-started-with-platform-encryption-for-data-cloud

Oct 2025 CDP Pulse Report - Salesforce Data 360

Yes/No	Details	Source(s)
14. Does your solution provide reverse-ETL / audience-syndication connectors that push profiles and events to major ad, analytics, and messaging destinations natively?	Yes. Salesforce Data 360 includes native audience activation connectors to push customer profiles and events to a wide range of destinations. Marketers can natively syndicate segments to major ad platforms like Google Ads, Facebook/Meta, LinkedIn, Amazon Ads, etc., directly from Data Cloud. Similarly, connectors exist to send data to email/SMS marketing systems (including Marketing Cloud Engagement) and analytics tools. Salesforce also partners with data onboarding services (LiveRamp, The Trade Desk, etc.) via AppExchange to extend reach. This "reverse ETL" capability is built-in – for example, a segment of high-value customers can be published to Facebook Ads or Google Customer Match with a few clicks, without custom coding.	https://www.salesforce.com/resources/articles/cdp-data-activation-guide-conf/ https://www.salesforce.com/resources/articles/cdp-data-activation-guide-conf/
15. Does your solution expose a unified customer-profile API/stream that any downstream system can query or subscribe to in real time?	Yes. Salesforce Data 360 provides a Unified Profile API that allows external systems to retrieve customer profile data on demand. Developers can use REST API endpoints (or GraphQL via the Salesforce API framework) to query unified profiles by ID or attributes in real time. This means downstream systems can call the API to get the latest 360° customer profile (including all unified fields) for a given customer and use it for personalization or analytics. Additionally, Data Cloud can push profile/event changes out via platform events or Data Actions for real-time subscriptions. For instance, if a profile is updated, a subscribing service could be notified immediately. This ensures any system can access or receive the up-to-date customer data from the CDP in real time.	https://developer.salesforce.com/docs/data/data-cloud-query-guide/references/data-cloud-query-api-reference/c360a-api-profile-dmname-id.html https://developer.salesforce.com/docs/data/data-cloud-query-guide/references/data-cloud-query-api-reference/c360a-api-profile-dmname.html
16. Does your solution publish transparent, usage-based pricing (events, MTU, storage) in writing, independent of professional-services spend?	Yes. Salesforce Data 360 uses a transparent consumption-based pricing model that is documented publicly. Pricing is based on usage measured in "Consumption Credits" for events, profiles, API calls, storage, etc., rather than fixed licenses. Salesforce provides a detailed rate sheet and pricing calculator outlining the credit costs for various operations (e.g. ingesting 1 million events, storing 1 TB of data). This information is available in writing on Salesforce's website. The pricing is separate from any services – you pay for the data volume and features you use, and professional services are optional. In short, Data Cloud's costs scale with usage and the pricing structure is made transparent to customers (allowing them to estimate and control spend).	https://www.salesforce.com/data/pricing/ https://www.salesforce.com/data/pricing/

Vendor Deep Dive



Adobe Experience Platform

Primary Archetype: Suite CDP

Strengths: Real-time/Edge; Governance & Residency

Capability Lens: Data Visibility & Quality;
Customization & Composability

VOX Score: 93.91/ 100

Preference Coverage: 15 / 16 questions “Yes”

📞 Talk to Adobe Experience Platform: <https://business.adobe.com>

📞 [Ask VOX for Introduction](#)

Oct 2025 CDP Pulse Report - Adobe Experience Platform

Yes/No	Details	Source(s)
Yes	Adobe Experience Platform (the foundation of Adobe's Real-Time CDP) includes built-in monitoring and observability tools. The platform offers a Monitoring Dashboard that provides near real-time visibility into data flows, showing the journey of data from ingestion sources through identity resolution, profile storage, segmentation, and activation destinations. These dashboards effectively provide end-to-end lineage for first-party data streams. Metrics and statuses for each stage (from source to profile to audience to destination) are updated continuously (typically within seconds), allowing users to quickly spot issues or delays. In short, Adobe's CDP gives users a live data observability interface to monitor pipelines and ensure data is flowing correctly across the system.	https://experienceleague.adobe.com/en/docs/experience-platform/dataflows/ui/monitor
Yes	Adobe Real-Time CDP provides a broad set of native source connectors and APIs to ingest data from many popular commerce, marketing automation, and customer service systems. Through Adobe Experience Platform's Sources catalog, users can easily connect to systems like e-commerce platforms (e.g. Adobe Commerce/Magento, Shopify, SAP Commerce), marketing automation tools (e.g. Adobe Marketo Engage, Salesforce Marketing Cloud, Oracle Eloqua, HubSpot, Mailchimp), and service/CRM platforms (e.g. ServiceNow, Salesforce CRM via APIs). These connectors are pre-built and do not require custom coding or professional services for standard integrations – they include guided UIs or REST API endpoints for authentication, scheduling data imports, and mapping fields. In summary, Adobe's CDP natively supports ingesting data from virtually all top commerce, email/SMS marketing, and support platforms out-of-the-box.	https://experienceleague.adobe.com/en/docs/experience-platform/sources/home https://experienceleague.adobe.com/en/docs/experience-platform/sources/home https://experienceleague.adobe.com/en/docs/experience-platform/sources/connectors/customer-success/servicenow
Yes	Adobe's data collection framework covers both client-side and server-side capture of customer events across all major channels. Using the Adobe Experience Platform Web SDK (for websites) and Mobile SDK (for iOS/Android apps), developers can collect behavioral data from web pages and mobile apps on the client side. These SDKs send data into Adobe's Edge Network in real time. For server-side and other sources, Adobe provides a Server-to-Server API (HTTP streaming endpoint) that can ingest events from backend systems – this is often used for importing email engagement events or point-of-sale transactions. All these collection methods feed into the same unified profile system. Essentially, Adobe's CDP lets you instrument once and unify data: whether events originate from a browser, a mobile app, an email service, or an in-store POS, the data can be streamed into the platform through a common pipeline, ensuring a single customer profile is updated in real time.	https://experienceleague.adobe.com/en/docs/experience-platform/sources/home https://hightouch.com/blog/what-is-adobe-cdp
Yes	Deterministic identity resolution is a core capability of Adobe's Real-Time CDP. The platform's Identity Service can ingest multiple identifiers (across online and offline touchpoints) and link them to the same person through defined identity namespaces. For example, Adobe's CDP will use a customer's login or loyalty ID, associated email/phone, device IDs (mobile Ad IDs like IDFA/GAID), and web cookies (ECID, etc.) and deterministically merge these into a single unified profile when matches occur. The system builds an identity graph for each individual using declared identifiers and any provided ID mappings (e.g. a POS transaction record with a loyalty ID will be attached to that profile). This ensures that behaviors from different channels – web visits, mobile app usage, in-store purchases – are stitched together as one customer. All identity linking in Adobe's CDP is deterministic (rule-based on ID matches), not probabilistic, so profiles are merged only when identifiers exactly match, giving a reliable single customer view.	https://experienceleague.adobe.com/en/docs/experience-platform/identity/home https://experienceleague.adobe.com/en/docs/experience-platform/identity/home
Yes	Adobe's Real-Time CDP is designed for fast ingestion and processing of streaming data, which covers use cases like inventory updates from online and physical stores. The platform supports streaming data ingestion, and once an event (such as an inventory count change) is ingested, it can update customer profiles or trigger segment re-evaluation within minutes. In fact, Adobe's documentation specifies that streaming data can qualify profiles for segments typically within a few minutes (usually under 5 minutes) of the event arriving. For even faster reaction, Adobe also offers an Edge Network that can process certain events in milliseconds (used for immediate personalization), but for inventory syncing specifically, the normal streaming pipeline is sufficient. In practice, this means that when a product is bought in-store or an inventory level changes, that information can be reflected in the CDP and made available for "Buy Online Pickup In Store" or endless aisle scenarios almost in real time (well within a 5-minute window 95% of the time).	https://helpx.adobe.com/legal/product-descriptions/real-time-customer-data-platform.html https://helpx.adobe.com/legal/product-descriptions/real-time-customer-data-platform.html
Yes	Yes. Adobe's experience stack includes real-time decisioning APIs that can respond within sub-second latency, enabling in-session personalization. Specifically, Adobe offers an Offer Decisioning service (part of Adobe Journey Optimizer/Experience Platform) which provides APIs to retrieve the "next-best offer" or personalized content for a given profile in real time. This service is built for high performance – it can evaluate a user's unified profile and return a decision typically in a few hundred milliseconds. (For example, Adobe's documentation notes that personalized offers from the decisioning engine load in about 500 ms on average.) Additionally, if using Adobe Target (Adobe's web personalization engine) in conjunction with Real-Time CDP segments, responses are similarly fast (Target's global edge network is optimized for <1s responses). In summary, Adobe's solution does expose APIs for real-time decisions/personalization, and they are engineered to respond quickly enough (usually well under one second) to allow actions to be taken while the user is still engaged on the site or app.	https://helpx.adobe.com/legal/product-descriptions/offer-decisioning-app-service.html

Oct 2025 CDP Pulse Report - Adobe Experience Platform

Yes/No	Details	Source(s)
7. Does your solution include embedded predictive-merchandising models (e.g., next-best product) deployable by business users without data-science help?	Yes Adobe Real-Time CDP includes embedded AI/ML capabilities geared for marketers, so they can leverage predictive insights without coding or data science expertise. In particular, Adobe provides Customer AI, an out-of-the-box machine learning service within the CDP that can generate propensity scores (for example, likelihood to convert, churn, or purchase a certain category). Business users can configure these models via a guided UI – defining the outcome of interest – and the system will train the model on historical data and output propensity scores for each customer profile. These scores (like “likelihood to buy product X”) are then attached to profiles and can be used directly in segmentation or personalization. All of this happens without the user needing to build models from scratch; the platform handles the heavy lifting. Additionally, Adobe’s integration of AI extends to features like look-alike modeling and product recommendations (often done through Adobe Target’s AI, a.k.a. Adobe Sensei). Overall, yes – marketers using Adobe’s CDP can deploy predictive merchandising use cases (like next-best product or propensity-to-buy models) through built-in AI services, no data scientist required.	https://experienceleague.adobe.com/en/docs/experience-platform/rtcdp/segmentation/customer-ai https://experienceleague.adobe.com/en/docs/experience-platform/rtcdp/segmentation/customer-ai
8. Does your solution allow teams to upload and deploy their own ML/AI models (PMML/ONNX/Python) to score profiles or events in real time?	Yes Adobe’s platform supports bring-your-own-model capabilities, though with some caveats. Adobe Experience Platform (on which the CDP runs) historically included a Data Science Workspace that allowed data science teams to develop custom models (in Python, using Jupyter notebooks) and then deploy these models to score data. Teams could import their own code or even package external algorithms as custom “recipes” which could be uploaded and run on the platform. Once a model is trained, it can be published as a scoring service and scheduled or invoked to score profiles/events, with the results written back into profiles for real-time use. It’s worth noting that these features typically operate in batch or near-real-time rather than millisecond in-line scoring; however, they do enable custom ML logic to be integrated. In summary, Adobe’s CDP environment is extensible: your team can bring custom ML models (developed in Python or exported through supported frameworks) into the platform, deploy them, and have their output incorporated into profiles or segmentation, using Adobe’s tooling and infrastructure.	https://experienceleague.adobe.com/en/docs/experience-platform/data-science-workspace/models-recipes/create-publish-model https://experienceleague.adobe.com/en/docs/experience-platform/data-science-workspace/models-recipes/create-publish-model
9. Does your solution provide an SQL workspace or direct lake-house access for segmentation, attribution, and ad-hoc analytics?	Yes Adobe Real-Time CDP provides an SQL query workspace called Query Service, which gives analysts direct access to the underlying data (the Experience Platform data lake) for advanced analysis. Users can write SQL queries (via a web UI or API) to join and analyze profile attributes, event data, experience events, etc. This is useful for building custom segments, performing attribution analysis, or any ad-hoc exploratory analytics beyond what the UI offers. The platform’s Query Service can query data in both the raw data lake and the profile datasets, enabling flexible analysis. Additionally, Adobe allows data export to external warehouses or BI tools if needed (there’s a dataset export feature), but many use cases can be accomplished directly in-platform with SQL. In short, marketers and data teams do have SQL capabilities within Adobe’s CDP to explore data and create custom insights, which complements the friendly UI segmentation with full analytical freedom when required.	https://helpx.adobe.com/legal/product-descriptions/real-time-customer-data-platform.html
10. Does your solution run continuous data-quality and anomaly detection with automatic quarantine or alerting before bad data feeds downstream systems?	Yes Adobe’s solution has multiple features to ensure data quality and catch anomalies before they cause issues downstream. First, the ingestion pipeline performs validation checks (type checking, schema validation, required fields, etc.) on incoming data; any records that fail validation are automatically diverted (quarantined) into a “bad data” holding area in the data lake rather than contaminating profiles – effectively preventing invalid data from flowing through to profiles or destinations. Additionally, Adobe Experience Platform provides Observability Insights and Alerts: users can set up alerts for various conditions (for example, if the volume of events drops or spikes unexpectedly, or if a dataflow fails). These alerts can notify teams via email or other channels so that anomalies in data ingestion or processing are caught in near real time. While Adobe may not call it a full “automated anomaly detection” system in the marketing sense, it offers the practical tools: automated data quality enforcement at ingest, and monitoring/alerting dashboards to detect and respond to irregularities.	https://experienceleague.adobe.com/en/docs/experience-platform/ingestion/quality/overview https://experienceleague.adobe.com/en/docs/experience-platform/observability/home
11. Does your solution enforce role-based access controls and immutable audit logs at the field level across all regions?	Yes Adobe Experience Platform (and thus Real-Time CDP) provides enterprise-grade security and audit features. Role-based access control (RBAC) is implemented, and notably Adobe supports attribute-based access control that can restrict access at the field-level. This means administrators can label sensitive data fields and grant permissions such that only certain roles can view or use those fields – ensuring compliance with privacy rules across regions. On the auditing side, Adobe maintains immutable audit logs of user and system activities. Every significant action (data ingestion, schema changes, segment publishes, permission changes, etc.) is recorded with who did it and when, forming a tamper-proof audit trail. These logs are accessible via the UI or API for review and are retained for compliance purposes. In summary, Adobe’s solution lets you tightly govern who can access what data (down to individual attributes) and keeps a comprehensive, unalterable log of all operations, which is crucial for security and regulatory audits in any region.	https://experienceleague.adobe.com/en/docs/experience-platform/sources/home https://experienceleague.adobe.com/en/docs/experience-platform/landing/governance-privacy-security/audit-logs/overview

Oct 2025 CDP Pulse Report - Adobe Experience Platform

Yes/No	Details	Source(s)
12. Does your solution offer built-in consent & preference management that suppresses activation instantly when permissions change (GDPR/CCPA ready)?	Yes Adobe's CDP has built-in consent and preference management capabilities to help comply with GDPR, CCPA, and similar regulations. It provides a standard data model for consent (so you can ingest consent signals – like a user opting out of marketing – from your consent management platform or website forms). Once consent data is in the platform, Adobe's Data Governance and Policy engine can automatically enforce it. For example, you can define policies that say "if a profile has withdrawn email marketing consent, exclude them from any email activation". These policies are evaluated in real time during audience activation. In practice, if a user's consent status changes (say they revoke permission), Adobe's system will immediately mark that profile and suppress them from any new audience exports to marketing destinations. When marketers try to activate an audience, the system filters out profiles that do not meet consent requirements on the fly, and it provides an explanation (e.g. reduced audience count due to consent). This ensures that activation is always in sync with the latest consent, fulfilling "do not target" requests virtually instantaneously. The platform is, therefore, GDPR/CCPA-ready by design – offering consent tracking and enforcement out-of-the-box.	https://experienceleague.adobe.com/en/docs/platform-learn/tutorials/data-governance/enforce-consent
13. Does your solution support regional data residency options (e.g., EU-only, US-only) with encryption at rest (AES-256) and in transit (TLS 1.2+)?	Yes Adobe Experience Cloud can accommodate regional data residency needs. Adobe has deployed data centers in multiple regions (North America, Europe, Asia Pacific, and even country-specific ones like a planned Canadian data center for AEP) – customers can have their Adobe Real-Time CDP instance and data hosted in a region that aligns with their compliance requirements (e.g., EU-only data storage to satisfy GDPR data localization). In terms of security, Adobe meets enterprise standards: all data at rest is encrypted (Adobe's platform uses cloud provider encryption, typically AES-256, for data in the Data Lake and profile stores) and all data in transit is encrypted via TLS 1.2+. The platform only communicates over HTTPS TLS 1.2 or higher for data ingress/egress, ensuring secure transport. Additionally, Adobe offers options for customer-managed encryption keys for those who need extra control. In summary, Adobe's CDP provides regional hosting options to keep data within required geographies, and it implements strong encryption for both data storage and data transit by default.	https://business.adobe.com/blog/the-latest/how-adobe-is-investing-in-canada-with-data-residency-and-helping-businesses-get-ahead-of-proposed-legislation https://experienceleague.adobe.com/en/docs/experience-platform/landing/governance-privacy-security/encryption
14. Does your solution provide reverse-ETL / audience-syndication connectors that push profiles and events to major ad, analytics, and messaging destinations natively?	Yes Yes. Adobe Real-Time CDP includes a Destinations framework which is essentially native reverse-ETL capability built into the platform. It has a catalog of pre-built destination connectors to all major marketing and advertising platforms – for example, you can natively push audience segments (with profile attributes) to Facebook Ads (as Custom Audiences), Google Ads and DV360, LinkedIn, Snapchat, Twitter, and other adtech platforms, as well as to email/SMS marketing systems (Adobe Campaign, Marketo Engage, third-party ESPs) and analytics tools. These connectors handle the data formatting and API integration required by the external systems. Additionally, there are "streaming destinations" and batch file destinations for data warehouses or cloud storage (like sending data to Amazon S3, Azure Blob, etc. for analytics). Adobe's CDP UI lets you simply choose a destination, map the fields, and the platform will continuously sync your chosen segments or event data to that target on a schedule or in real-time. No custom coding is needed – the integrations are out-of-the-box. In short, Adobe provides robust audience syndication capabilities to push profile data and event streams to virtually any important downstream system in the marketing stack.	https://experienceleague.adobe.com/en/docs/experience-platform/destinations/home https://experienceleague.adobe.com/en/docs/experience-platform/destinations/home
15. Does your solution expose a unified customer-profile API/stream that any downstream system can query or subscribe to in real time?	Yes Adobe's Real-Time CDP exposes a Unified Profile API that allows external systems to query customer profiles on demand. Developers can use RESTful calls to the Experience Platform's Real-Time Customer Profile endpoint to retrieve a profile by an ID (e.g., get all attributes and segment memberships for a given user). This API gives programmatic access to the single customer view, so any authorized downstream application can look up the latest profile data in real time. Moreover, Adobe supports streaming export of profiles and events: with the Advanced Event Forwarding/Stream destinations (available in certain editions), the CDP can publish profile updates or audience qualification events in real time to message queues or webhooks (for example, to an AWS Kinesis stream, Azure Event Hub, or a custom HTTP endpoint). This means systems can also subscribe to those streams to get incremental changes. In summary, Adobe provides both a queryable profile API and options for real-time subscription to data streams, enabling other applications to leverage the unified profile data in real time.	https://experienceleague.adobe.com/en/docs/experience-platform/profile/api/entities https://experienceleague.adobe.com/en/docs/experience-platform/destinations/destination-types
16. Does your solution publish transparent, usage-based pricing (events, MTU, storage) in writing, independent of professional-services spend?	No Adobe's pricing for Real-Time CDP is not published transparently in a usage-based rate card. Adobe typically uses a custom pricing model negotiated with each client. The costs are primarily based on key usage metrics like the number of customer profiles under management (and possibly volumes of events or data storage), but the exact prices for tiers are not publicly posted on their website. Instead of a simple usage fee list, Adobe offers various editions (e.g., Standard, Prime, Ultimate) and one must engage with Adobe's sales team to get a quote. In other words, you won't find a publicly available document listing "\$X per 1000 events" or "\$Y per MAU" – pricing is provided via sales and often bundled with other Experience Cloud products. It is independent of any professional services in the sense that software licensing is separate, but the lack of publicly transparent pricing means cost estimates require consultation with Adobe.	https://hightouch.com/blog/what-is-adobe-cdp

Vendor Deep Dive



ReactorData

Primary Archetype: Suite CDP

Strengths: Governance & Residency

Capability Lens: Data Visibility & Quality; Security,
Consent & Residency; Customization & Composability

VOX Score: 69.01/ 100

Preference Coverage: 11 / 16 questions “Yes”

📞 Talk to ReactorData: <https://reactordata.com>

📞 [Ask VOX for Introduction](#)

Oct 2025 CDP Pulse Report - ReactorData

	Yes/No	Details	Source(s)
1. Does your solution provide data-observability dashboards that refresh in ≤ 60 seconds (p95) and show source-to-destination lineage for every first-party data stream?	Yes	ReactorData includes a real-time Data Monitor dashboard for tracking data pipelines. Users can monitor each data source's status, last refresh, and health in one place, with updates in near-real time. The platform emphasizes end-to-end data visibility – you can trace and validate data transformations from sources through to destinations, promoting transparency and trust in every first-party data stream.	https://docs.reactordata.com/hc/en-us/articles/37275217616276-Data-Monitor-Overview https://soundcommerce.com/solutions/retail-operations/ https://soundcommerce.com/why-soundcommerce/soundcommerce-vs-cdps/20 https://soundcommerce.com/solutions/retail-operations/
2. Does your solution ship native inbound connectors or APIs for the top-10 commerce, marketing-automation, and service platforms without custom code or PS?	Yes	ReactorData (formerly SoundCommerce) provides a broad library of pre-built connectors covering popular commerce, marketing automation, and service platforms. It natively integrates with hundreds of retail data sources – for example, e-commerce platforms (Shopify, Adobe/Magento, BigCommerce, Salesforce Data 360, etc.), email/SMS marketing tools, digital ad platforms, web analytics, POS systems, and more – all without custom coding. This allows data ingestion from the top commerce and marketing systems out-of-the-box	https://soundcommerce.com/solutions/it-data/ https://soundcommerce.com/resources/integrations-and-partners/ https://soundcommerce.com/solutions/it-data/ https://soundcommerce.com/resources/integrations-and-partners/
3. Does your solution support server-side and client-side collection for web, mobile-app, email, and in-store (POS/scanner) events in a single SDK?	No	ReactorData focuses on server-side data integration rather than a unified client SDK for event collection. It does not offer its own client-side web or mobile SDK. Instead, ReactorData ingests data from source systems (e.g. pulling e-commerce orders, POS transactions, email campaign data via APIs) and can incorporate third-party behavioral streams. For web/app events, it relies on integrating data from tools like Google Analytics, Snowplow, or Segment rather than direct on-page/tag SDK collection. In-store and POS events are integrated via connectors to those systems, not through a single client-side SDK.	https://soundcommerce.com/platform/platform-faqs/ https://soundcommerce.com/platform/platform-faqs/ https://soundcommerce.com/platform/platform-faqs/ https://soundcommerce.com/platform/platform-faqs/
4. Does your solution offer deterministic identity-resolution that stitches online cookies, MAIDs, loyalty IDs and POS receipts into one customer profile?	Yes	ReactorData implements deterministic omnichannel identity resolution. SoundCommerce's platform creates a Customer Unique ID (CUID) to unify profiles, using direct matching and deduplication across identifiers. It can stitch together online and offline customer identifiers – for example, linking website visitor IDs (or cookies and MAIDs) with e-commerce customer accounts, loyalty program IDs, and in-store transaction records (receipts) – into a single customer profile. This matching is deterministic (rule-based) to ensure one unified profile per customer.	https://soundcommerce.com/platform/platform-faqs/ https://soundcommerce.com/platform/features-benefits/ https://soundcommerce.com/platform/platform-faqs/ https://soundcommerce.com/platform/features-benefits/
5. Does your solution sync online and in-store inventory events in < 5 minutes (p95) to enable BOPIS and endless-aisle use cases?	Yes	ReactorData's pipeline is designed for low-latency data updates, enabling near real-time synchronization of inventory data across channels. It can ingest inventory changes from online stores and in-store POS systems rapidly (on the order of minutes). This timely integration of inventory events supports use cases like "Buy Online, Pickup In Store" (BOPIS) and endless aisle, since product availability updates in one channel are quickly reflected in the unified data model for other channels. The platform connects product catalog and inventory data from all sources to provide a unified, up-to-date view of stock levels by location. (While specific SLA like p95 < 5 min isn't publicly stated, the system emphasizes real-time data integration for inventory and orders.)	https://soundcommerce.com/why-soundcommerce/soundcommerce-vs-cdps/ https://soundcommerce.com/solutions/retail-operations/ https://soundcommerce.com/why-soundcommerce/soundcommerce-vs-cdps/ https://soundcommerce.com/solutions/retail-operations/
6. Does your solution expose decision or personalization APIs that return in < 1 second (p95) so actions fire while the session is active?	No	ReactorData does not provide a native real-time personalization or decisioning API that guarantees sub-second responses. The platform's strength lies in preparing and modeling data for analytics and activation, rather than delivering on-the-fly web personalization decisions within active user sessions. In practice, real-time decisioning (e.g. on-site product recommendations or personalized content within 1s) would be handled by downstream customer engagement or personalization tools integrated with ReactorData. For instance, SoundCommerce feeds data to marketing platforms like Cordial or Braze for activation, instead of an in-house API responding directly to live session calls.	https://soundcommerce.com/resources/whats-new/ https://soundcommerce.com/resources/whats-new/
7. Does your solution include embedded predictive-merchandising models (e.g., next-best product) deployable by business users without data-science help?	Yes	ReactorData offers built-in predictive analytics models that business users can leverage. In particular, the platform includes machine-learning driven models for customer analytics – for example, forecasted Customer Lifetime Value (CLV) and repeat purchase propensity scoring are provided out-of-the-box. These predictive models (while focused on customer value and behavior) are embedded in the product and accessible via a no-code interface, so non-technical users can use them for segmentation and decisioning. (Specific "next-best product" recommendations aren't explicitly mentioned in the documentation, but the presence of predictive CLV and propensity models indicates similar predictive merchandising insights are available to users without needing a data scientist.)	https://soundcommerce.com/why-soundcommerce/soundcommerce-vs-cdps/ https://soundcommerce.com/platform/features-benefits/ https://soundcommerce.com/why-soundcommerce/soundcommerce-vs-cdps/ https://soundcommerce.com/platform/features-benefits/ https://soundcommerce.com/why-soundcommerce/soundcommerce-vs-cdps/

Oct 2025 CDP Pulse Report - ReactorData

	Yes/No	Details	Source(s)
8. Does your solution allow teams to upload and deploy their own ML/AI models (PMML/ONNX/Python) to score profiles or events in real time?	No	<p>There is no indication that ReactorData (SoundCommerce) supports uploading custom ML or AI models for deployment within the platform. The solution does not advertise a feature to import user-provided models (e.g. PMML/ONNX files or custom Python code) for real-time scoring on incoming data. Instead, ReactorData focuses on its built-in data models and works alongside external AI/analytics tools. Users can always export data to their own data science environments or use the data in a warehouse to apply custom models outside ReactorData, but the platform itself doesn't serve as a real-time model hosting service for arbitrary user models.</p>	https://soundcommerce.com/why-soundcommerce/soundcommerce-vs-cdps/ https://soundcommerce.com/why-soundcommerce/soundcommerce-vs-cdps/
9. Does your solution provide an SQL workspace or direct lake-house access for segmentation, attribution, and ad-hoc analytics?	Yes	<p>ReactorData is built on a "composable" data stack paradigm, which means customers retain direct access to the data in their cloud warehouse. SoundCommerce lands the unified data in the client's Snowflake or BigQuery (which the client owns and controls), enabling analysts to run SQL queries directly on that data. The platform also includes a lightweight no-code query workspace for segmentation and analysis, but importantly, advanced users can use SQL or connect their preferred BI tools to the underlying data models for ad-hoc analytics and attribution exercises. This ensures full flexibility for custom queries and analysis on the CDP's data.</p>	https://soundcommerce.com/why-soundcommerce/soundcommerce-vs-cdps/ https://soundcommerce.com/why-soundcommerce/soundcommerce-vs-cdps/ https://soundcommerce.com/platform-faqs/
10. Does your solution run continuous data-quality and anomaly detection with automatic quarantine or alerting before bad data feeds downstream systems?	Yes	<p>ReactorData places a strong emphasis on data quality and reliability in its pipelines. The platform uses Data Standards (predefined schema and validation rules) for each source to ensure incoming data meets consistency and quality requirements. Data is validated at ingestion – if it doesn't conform, it can be flagged or stopped before corrupting downstream models. Additionally, the system provides monitoring of data connectors and can alert users to any issues (e.g. a paused or failing data feed). While the documentation doesn't explicitly mention an "auto-quarantine" feature, the combination of schema enforcement and proactive monitoring/alerting helps catch anomalies or bad data early so they can be addressed prior to affecting any outputs.</p>	https://docs.reactordata.com/hc/en-us/articles/37128827036308-Source-Links https://docs.reactordata.com/hc/en-us/articles/37275217616276-Data-Monitor-Overview https://docs.reactordata.com/hc/en-us/articles/37128827036308-Source-Links https://soundcommerce.com/solutions/retail-operations/
11. Does your solution enforce role-based access controls and immutable audit logs at the field level across all regions?	Yes	<p>ReactorData includes robust security and access control features. Administrators can define role-based access controls (RBAC) with the principle of least privilege – meaning users only see and manage the data they're permitted to. Permissions can be set at granular levels to align with company policies. The platform (SOC 2 Type 2 certified) also maintains extensive logging of system and user activities, providing an audit trail of data access and changes. These security measures apply across all deployments (and regions), ensuring compliance and traceability down to fine-grained levels (though field-level permission granularity is implied by the ability to control access to various data objects in the platform).</p>	https://soundcommerce.com/platform/security-compliance/ https://soundcommerce.com/platform/security-compliance/ https://soundcommerce.com/platform/security-compliance/ https://soundcommerce.com/platform/security-compliance/
12. Does your solution offer built-in consent & preference management that suppresses activation instantly when permissions change (GDPR/CCPA ready)?	No	<p>ReactorData does not have a native consent or preference management module for end-user privacy preferences. There is no built-in interface for capturing GDPR/CCPA consent statuses or automatically suppressing data activation based on consent changes. Instead, it assumes the business has obtained necessary consents for any data ingested. In fact, the SoundCommerce terms of service specify that the customer is responsible for having all required user consents for the data provided to the platform. While ReactorData is built with data security and privacy best practices (and can integrate with consent data if provided), it does not itself manage or update user opt-ins/opt-outs in real time as a dedicated consent management tool would.</p>	https://soundcommerce.com/saas-agreement/ https://soundcommerce.com/saas-agreement/
13. Does your solution support regional data residency options (e.g., EU-only, US-only) with encryption at rest (AES-256) and in transit (TLS 1.2+)?	Yes	<p>ReactorData can accommodate data residency requirements by leveraging the regional infrastructure of its cloud platform (it runs on Google Cloud, which offers multiple regional data centers). Clients with EU-only or US-only data needs can deploy ReactorData in those regions to keep data local. On security, SoundCommerce/Reactor uses strong encryption for data both in transit and at rest. All data transferred is protected via TLS, and data at rest on their cloud servers is encrypted (using robust standards like AES-256). The company highlights that it employs robust encryption and follows industry best practices to secure PII and business data.</p>	https://soundcommerce.com/platform/security-compliance/ https://soundcommerce.com/platform/security-compliance/
14. Does your solution provide reverse-ETL / audience-syndication connectors that push profiles and events to major ad, analytics, and messaging destinations natively?	Yes	<p>ReactorData supports activating customer data to external systems through integrated reverse-ETL capabilities. Rather than building its own outbound connectors from scratch, SoundCommerce partnered with leading reverse-ETL tools (e.g. Census is embedded in the platform) to natively push audiences and events to popular destinations. This means users can seamlessly sync customer profiles, segments, or event data from ReactorData into major ad networks (Google, Facebook, etc.), analytics platforms, CRM and customer engagement tools (like Braze, Cordial), and other marketing endpoints without custom code. The integration with these audience syndication connectors is handled within the UI, allowing data to flow out to downstream systems for activation as a built-in feature.</p>	https://soundcommerce.com/resources/whats-new/ https://soundcommerce.com/platform/platform-faqs/ https://soundcommerce.com/resources/whats-new/ https://soundcommerce.com/platform/platform-faqs/ https://soundcommerce.com/resources/whats-new/

Oct 2025 CDP Pulse Report - ReactorData

Yes/No	Details	Source(s)	
15. Does your solution expose a unified customer-profile API/stream that any downstream system can query or subscribe to in real time?	No	ReactorData does not expose a dedicated real-time "profile API" or streaming subscription feed for customer profiles. The unified customer profiles are accessible via the data warehouse or via batch connectors, but there isn't a published REST/GraphQL streaming API for on-demand querying of profiles in sub-second timeframes. Downstream systems get ReactorData's customer data through scheduled syncs or reverse-ETL integrations rather than an open subscription stream. For real-time needs, ReactorData relies on integration with realtime-capable marketing platforms (for example, in a case study SoundCommerce data feeds into Cordial to drive real-time messages). In summary, external systems cannot directly subscribe to a live profile/event stream from ReactorData; instead data is shared through the integration/activation layer.	https://soundcommerce.com/resources/w_hats-new/ https://soundcommerce.com/resources/w_hats-new/ https://soundcommerce.com/platform/platform-faqs/
16. Does your solution publish transparent, usage-based pricing (events, MTU, storage) in writing, independent of professional-services spend?	Yes	ReactorData provides clear, publicly documented pricing on its website. The pricing is usage-based and transparent: for example, as of current information, the "Core" plan is \$2,000/month for up to 30 million rows/events processed per month, and a "Pro" plan at \$4,000/month for up to 75 million rows, etc., with custom enterprise plans for higher volumes. Additional source connectors beyond the included ones are priced at a fixed rate (e.g. \$100 per extra source). These prices are published in writing and are not contingent on professional services engagements. This usage-based model (based on data volume and connectors) is clearly delineated, allowing customers to predict costs without hidden fees.	https://reactordata.com/pricing/ https://reactordata.com/pricing/ https://reactordata.com/pricing/ https://reactordata.com/pricing/



Empowering B2B Markets with Buyer Voice

VOX is a human-led, AI-native market research firm specializing in B2B markets. By dynamically capturing the preferences and priorities of senior executives in retail and eCommerce, VOX leverages proprietary AI throughout its research process to objectively amplify the voice of leaders.



Our mission is to align both sides of the market using real-time insights directly sourced from industry leaders. VOX provides unique and actionable insights to both senior executives and vendors, helping organizations effectively align their go-to-market strategies with the evolving priorities and preferences of key decision-makers.



FOR LEADERS

Executive Peer Intelligence

Gain actionable insights, peer benchmarking, and monetary rewards by sharing your preferences and priorities.

[Join Vox Leaders](#)



FOR VENDORS

Real-Time Market Insights

Align marketing and sales strategies with real-time buyer voice, achieving predictable pipeline growth.

[Discover the Buyer Voice Platform](#)



www.voxwins.com