blueconic

BLUECONIC PRODUCT PROFILE

What is the product name?
BlueConic Customer Data Platform

What problem does it solve?
To deliver an end-to-end customer experience, modern organizations realize that an array of functions and teams are required, including marketing, commerce, analytics, customer support, and more. However, the various tools of their trades (email systems, campaign management platforms, analytics & BI tools, DMPs, ad platforms, ecommerce platforms, POS, CRM) are siloed within each department and thus, store data and recognize customers/prospects in ways that are unique to each system.

What these teams and tools lack is access to first-party customer data that is both unified and actionable, which is required to effectively orchestrate individualized/bespoke experiences across all touchpoints with their customers. As a consequence, departments responsible for driving the growth of the company are at the mercy of other entities – both internal (e.g., IT, the analytics team, or legacy technologies/databases) and external (e.g., agencies, technology vendors) – to dictate what, when, where, and how they can use data to orchestrate customer experiences.

BlueConic not only unifies first-party customer data at an individual level, but equally important, it also delivers that unified profile data back to the business technology users in marketing, ecommerce, analytics, etc. in a format that their tools can use to improve how they engage with customers, conduct modeling and analytics, build segments, and more.

What makes it unique/sets the product apart apart/ how it excels above the competition?
BlueConic can _flex_ in all the ways a modern organization needs it to – namely in its speed, adaptability, value, and utility to business technology users. That's because everything in BlueConic starts with the unified profile.

Unlike other CDPs, every data point collected in BlueConic is translated into an attribute of an individual customer that is persistently stored in the profile, making it possible to quickly and confidently query the data in support of a wide range of activation use cases, including lifecycle marketing orchestration, modeling & analytics, digital products and experiences, and/or audience-based monetization.

Whether an organization is at the beginning of its digital transformation or looking to take its transformation to the next level, BlueConic provides the means to handle the volume, variety, and velocity of data that will be required power it. In other words, BlueConic scales for today's goals and tomorrow's ambitions.

Data Collection – How robust is the data that the CDP can natively collect? How difficult is it to capture custom audience/behavioral data?
BlueConic provides an extensive set of point-and-click data collectors (or "listeners") to customers to natively collect data onsite. We also offer out-of-the-box behavioral scores (engagement, intensity, recency, etc.). Additionally, we have a number of named and universal connections to other online and offline technologies which allows for unlimited data collection capabilities with BlueConic. All our connections have standard profile property mapping models that can be easily configured by users using a web-based UI.

Data Latency – how quickly does the data get manipulated within the CDP?

BlueConic's customer data platform is real time, and therefore almost all data ingestion and merging happens on the fly. Certain systems, however, only allow for batch exporting, and therefore, will be batch imported to BlueConic. As soon as data changes in a profile, they're processed immediately, including profile merges, which occur instantly, in under 24 milliseconds. Any changes to the profile that are tied to segment rules also happen immediately, so when a profile changes, segments are updated as relevant. Therefore, any activation associated with that segment (e.g., a Lifecycle stage, a connection, a personalized experience, and/or a report) will all be updated dynamically as well.

Data Science - in what ways does the CDP incorporate data science concepts like machine learning and automatic optimization?

BlueConic provides behavioral scoring to inform segmentation and activation of customer data as well as a recommendations engine to provide product/content recommendations online and in email. In 2019, we introduced AI Workbench, which enables data scientists to import or build models within BlueConic using Jupyter notebooks. It also enables marketers and other business users without a data science background to select out-of-the-box models from our gallery (e.g., propensity, customer lifetime value, RFM, lookalike, etc.).

Because the models exist within the same system as the customer's unified profile database, users can immediately segment and activate on the output of those models using our certified connections. BlueConic also offers profile grouping features to make it easier for customers to group related profiles into households, accounts/companies, or other types of groups their business cares about.

Ease of Use - Can the user easily achieve their desired results? Is the UI intuitive and easy to navigate? Is the platform easy to learn for tech and non-tech users alike?

Our UI is often praised by customers and analysts alike. While some CDPs purely focus on analytics and personalization, and others focus heavily on operational data management without any UI for segmentation and activation. BlueConic is unique in the market in that our UI enables both business users and technical users to use the platform to suit their needs. For example, AI Workbench provides a data scientist with the ability to create or import Jupyter notebooks with their company's machine learning models to BlueConic, while a marketer has the ability to schedule these notebooks to run so they can refresh customer scores for segmentation as often as needed. Marketers can also use the point-and-click UI to test how different parameters – like a specific geolocation – influences an audiences' propensity to buy without having to ask data scientists to build a brand-new model.

Export and Import - how well does the CDP import and export data out of its platform? How many methodologies are available to the marketer to do so?

BlueConic provides more than 80 certified, out-of-the-box connections with other products and technologies (a full list can be found at www.blueconic.com/connections). There are also a number of universal methods of connection that we employ to import and export data. Data import and exports can happen in real-time or can be scheduled as a batch run. Data imported in BlueConic is also processed in real-time, meaning all data transformation happens instantaneously. Depending on a specific customer's need for importing and exporting data, there are seven ways to connect with BlueConic: https://www.blueconic.com/blog/customer-data-integrations/

Innovation – Is the product using cutting edge technology, strategies, ideas or features?

- <u>First-Party Hostname Console</u>: In February 2021, BlueConic introduced new functionality that enables customers to continue building legitimate first-party data assets amid rising third-party cookie and privacy-related data restrictions. Using a centralized, easy-to-use console, BlueConic customers can centrally manage the implementation of first-party server-side cookies to continue to build their first-party data assets and untether their marketing tactics from technologies that rely on third-party cookies. As a result, companies can continue capturing on-site data to enrich the unified customer profiles stored within their BlueConic database with confidence and utility, while insulating themselves from future browser privacy changes.
- <u>Lifecycle Orchestration enhancements</u>: Last summer, BlueConic introduced enhancements to its best-in-class lifecycle orchestration capabilities that enable marketers to use real-time data to optimize their programs based on where each individual customer or prospect is on their unique journey. The fully configurable capability frees marketers from rigid, rules-based workflows by enabling users to organize, coordinate, and visualize the marketing touchpoints used to influence and advance individuals from one customer lifecycle stage to the next.
- <u>Aggregated Data Manager</u>: To provide multi-brand companies with the ability to individually manage brand tenants and gain access to a cross-brand view of an individual, BlueConic is introducing Aggregated Data Manager. This functionality will give each brand team control over their first-party data, while also giving corporate teams the ability to combine attributes of individuals from across brands for cross-brand acquisition opportunities and advanced analytics.
- <u>Focus on the growth-focused buyer</u>: To date, CDPs have primarily been considered a tool for marketers who need to achieve a single customer view so they can orchestrate individualized experiences across channels and customer lifecycle stages. But to deliver a true, end-to-end customer experience, modern organizations are realizing that an array of functions and teams are required, including marketing, commerce, analytics, customer support, and more. As such, we've expanded our addressable market beyond marketing to all departments responsible for driving the growth of the company.

Integration - How well does the CDP support/integrate third party tools? How robust is the customer support in implementing the CDP?
As previously mentioned, BlueConic provides more than 80 certified, out-of-the-box connections with other products and technologies (a full list can be found at www.blueconic.com/connections), along with seven other ways to connect with BlueConic: https://www.blueconic.com/blog/customer-data-integrations/

BlueConic's Customer Success Framework assists customers in leveraging their BlueConic investment by providing expertise and content throughout the customer lifecycle. The framework distinguishes between onboarding activities (e.g., project management, use case identification and prioritization, training, etc.) and implementation projects for situations that require deeper discovery related to custom configurations, requirements, or data onboarding mechanisms in order to define a larger implementation project either with a fixed price/fixed scope or based on a time and material engagement.

After the initial onboarding or implementation period, BlueConic provides ongoing user support and supplemental training to accommodate continuous learning and to account for staff growth and churn within a team. We also provide ongoing consulting and strategic services. These are delivered either in a packaged format driven towards the successful implementation of a set of specific use cases via

access to content and training, or in the form of access to our team of skilled Data Scientists, Solution Architects, and Consultants.

Persistent IDs - How well does the system associate old and new data for a customer/visitor when there is a change in the associated identifier (email address, phone number, account ID, etc.)? How well does the CDP create and maintain a persistent ID?
BlueConic's profiling mechanism is at the core of all the other features in the platform. BlueConic profiles are unique in a few key ways:

1) the ability to create them for both known and unknown individuals and merge them based on unique identifiers
2) they exist across any channel and persist as long as the marketer requires
3) they are designed to pull from any other technologies in your stack – a CRM or email system, for example

Any identifier or combination of identifiers can be used for profile merging. BlueConic defines a unique identifier as an attribute that can be used to retrieve profiles with a specific value for a property, for example, an email address or a customer ID. Since the profile persists server side, the match between the two devices persists and will 'remember' the link between the two devices once the merge has been established.

Most BlueConic customers use deterministic profile merging methods, so there must be (at least one) verifiable data point of information shared by two profiles to merge them together. This differentiates a CDP from other platforms that rely on probabilistic identity stitching methods. The merger can be based any profile value or a combination of profile values. BlueConic detects when a customer switches from a channel and uses a one-time identifier to connect the ID on the new channel to the old – for example when a customer switches from an email to web, from web to another website, from a push message to an app, from email to app, etc. When this mechanism has not been triggered and an ID is found later in the journey, the database is searched and the profiles and all their historical data are merged.

BlueConic also provide probabilistic matching, or "fuzzy" matching, through an AI Workbench model that serves as a starting point. Customers can modify both via the UI and via the underlying Python code. This model uses a statistical algorithm and weights to determine match scores and thresholds to determine a match, which includes the Levenshtein Distance to measure similarity between words/names, weighted with some measure of entropy/word length, and approximate k-Nearest Neighbor (kNN) graphs, to get an online nearest neighbor algorithm.

Reporting and analytics – How robust are the native tools for reporting results?
BlueConic offers an unlimited number of easily configurable dashboards that bring all different types of data and insights together in clear, visual overviews within a point-and-click UI. Dashboards are comprised of BlueConic Insights and features include:

- Shared dashboards with other users
- Functionality to design and edit different types of visualizations and data calculations using a pluggable insights gallery
- Custom visualizations
- Dragging and dropping insights across the dashboard

- Designation of specific targets and reporting periods that enable users to instantly monitor the results against the target

Dashboards are fully customizable to the BlueConic user. By design, the user interface eliminates the need to query the data with SQL, though users are still able to set up queries to build segments. Dashboards can be built to report on segments, profile properties, and user interactions within the BlueConic platform. Moreover, data is also available as described, which enables analysis across a number of data visualization and statistical computing packages. Data is exportable directly into CSV or made accessible to other applications through out-of-the-box connections with the most common apps or with a custom connection via our REST APIs.

All Insights are highly configurable using our web-based UI. In addition to detailed reports that can be built within the platform, our extensive list of connections allows BlueConic users to be able to integrate with many of the data visualization and reporting tools that they are already using. For example, via the connection to Google's Cloud Platform (BigQuery), users can integrate BlueConic data with data from Adwords (or DoubleClick), Google Analytics, etc. in Google Data Studio.

For additional reports and analysis, BlueConic customers can also leverage AI Workbench, which offers the ability to run python code and models within the BlueConic system. Visualizations are generated within Jupyter notebooks and insights can pull from these models to show any data visualizations.

Scale of Data - What is the scale of data that can be ingested/ orchestrated on an ongoing basis?
BlueConic's architecture is built on top of a profile database (Apache Cassandra), which is a linearly scalable, fault-tolerant, and capable of storing billions of key/value pairs, and hundreds of terabytes of data. All profiles are updated in real time. Data is available immediately upon capture/ingestion. The association and consolidation of behavioral events into persistent, person-level profiles happens instantly, as does the identification of any segment criteria a profile matches; that process happens in real time on the server and has no effect on a site's responsiveness.

By not relying solely on an event stream database to associate identities, we allow customers to balance flexibility/intention in data collection and consolidation. BlueConic composes an actual "persistent" profile, server-side and in real time instead of creating a graph of associated, deconstructed events. The profile is stored in a NoSQL database with both high-volume storage capability and incredibly fast read/write speeds so that data is unified and actionable in specific ways. It's immediately available for segmentation, personalization, and analytics – all while preventing the cost of data retention of raw events that explode in the long term.

BlueConic's hybrid method of native big data profile storage, event/timeline processing, and the ability to deploy relational groupings (like households and accounts) offers marketers – via a UI – an easy way to persist and structure attributes (or to create those profile properties upon data import) and to map those attributes between systems in a few clicks, while ensuring the data remains available in real time for both read and write purposes at scale (in short, as is required for real-time marketing). With the grouping capability, users can create and manage Groups (like households, accounts, buildings, business entities, etc.) to organize and segment profiles. Groups can have group properties that can be managed using the UI or API as well. Bulk APIs are available for exporting and importing group data. Groupings are integrated in the BlueConic UI (such as search, navigation, pluggability, activity stream).