

# Real-Time Decisioning Puts the Customer in Charge of the Journey



**One of the biggest conundrums that businesses face today is optimally communicating with customers. It's tempting to target top customers nonstop, for example, with ads, emails, and offers about *all* of your products or services. And it can be worrying to omit customers from campaigns when there's even the slightest chance that they'll purchase; plus, limiting messages can inadvertently omit high-priority customers. But overcommunicating can lead to poor customer experiences and even churn. Nearly two-thirds of customers (59 percent) opt out of a brand's email communications because they receive too many emails, for example.**

So, how do you arbitrate your campaigns and deliver more personalized experiences? Taking an *a priori*, or presumptive approach to managing and optimizing communications rarely works and is complicated. A far better way to approach customer interactions is real-time decisioning. It's a simple and elegant solution that enables you to deal with all the complexity of omnichannel communications. Instead of having to be a traffic cop who needs to spend innumerable hours determining what car goes where and when, you can focus on strategy while real-time decisioning serves as perfectly timed traffic lights.

Real-time decisioning eliminates the traffic jams that prioritization can cause. It helps you determine what each customer needs to hear from you when—in any channel—creating always-on personalization rather than random acts of personalization that present as inconsistent experiences. Brands are now able act *a posteriori*, reacting in real time to observable facts

In other words, it eliminates the need to plan time and place in advance and puts customers at the center of their own journey—allowing you to reach them with a next-best action or offer in whatever channel will work most effectively at the time that's optimal for the customer. It's a channel-agnostic approach, ensuring that each customer will hear the most important thing *that customer* needs to hear at that moment, decoupled from "forced" channel interactions.

You can reduce the number of generic campaigns blasted at everyone and instead use rules to programmatically develop and deliver personalized messages that take into account a customer's previous behaviors and current context. It's a "derived by fact" approach that enables you to filter outbound messages and deliver the content that's the most relevant at

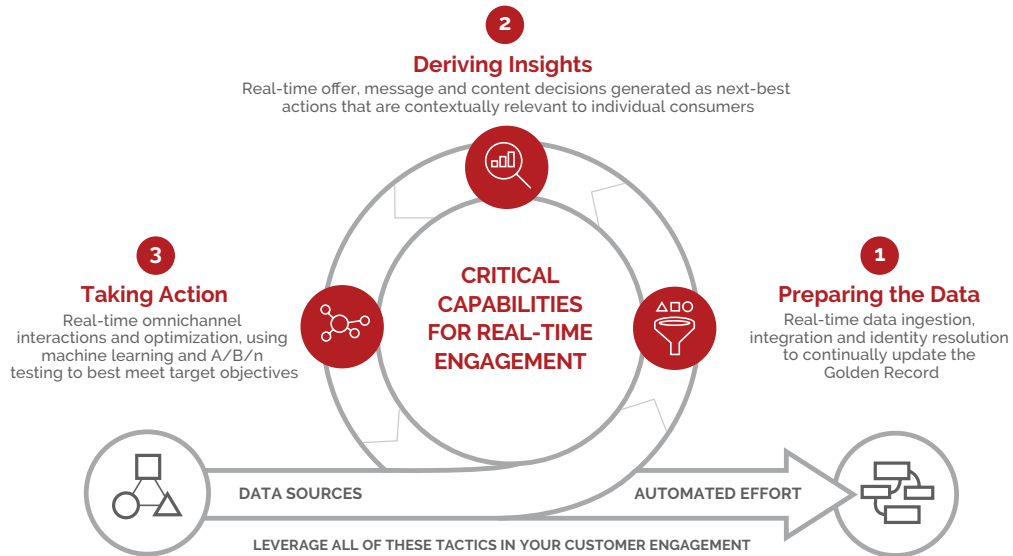
the time—including at the moment of open for email messages. In fact, using real-time decisioning with a dynamic email platform will ensure that the customer is always getting the best message timed up to the last second that they open it, even if they open an email days after they received it.

Using real-time decisioning to guide your customer communications is an innovative way of overcoming some traditional interaction-related challenges that cause friction and frustration, including message redundancy. It's a unique way of using customer data to drive a precise set of messages that you can automatically recalculate based on each customer's individual interactions and the overall volume of customer interactions. Also, you can automatically freshen your models based on any kind of trigger, without people needing to be involved.

## Quality data underlies decisioning

This is simpler and more effective way of personalizing customer interactions than a complicated presumptive approach, but it will only work with high-quality data that the real-time decisioning engine inside a customer data platform (CDP) can access. It requires a CDP with technology such as the Redpoint rgOne platform to provide a precise data environment with processes that structure and maintain customer data in real time to ensure that you'll always reach the right audience.

An option such as using a data lake and applying a data model at the moment of a campaign isn't as effective as a real-time decisioning using a CDP; nor is it safe enough to meet today's strict regulations. For instance, the former provides no traceability or data lineage. You can tell what you sent but you lose the "why" of the event that enables you to demonstrate that you sent something for reasons that are appropriate.



Ingesting and preparing data, deriving relevant insights, and taking action are the three main steps involved in setting up and using real-time decisioning for message-list-driven real-time personalization. They're also the three main steps to creating the closest approximation to that perfect sequence of messages, eliminating redundancy, ensuring relevance, and doing it all with the least amount of friction possible. Let's take a closer look.



### Step 1: Ingesting and Preparing the Data

Data quality is of the utmost importance for real-time personalization and decisioning. Data errors create experience errors, and consumers' thresholds for errors are low and getting lower. But, when it comes to data ingestion and preparation, little attention has been paid to the accuracy, correctness, and timeliness of the data being ingested. Much more investment has gone into the whizbang technologies that surround the execution channels than into tools that improve and simplify ingesting and preparing data.

One reason for this is the difficulty of getting data ingested, cleansed, structured, formatted, completed, and then matched together into a comprehensive, 360-degree view of a customer (i.e., a Golden Record). It takes extensive technical knowledge, considerable experience in matching, and the willingness to take on a set of processes that are never done. Today, however, an AI-based customer engagement platform such as rgOne can automate many of these processes.

There are many methods for moving or communicating data. These can vary from the traditional dropping of a file in a folder or using an FTP landing zone spanning all the way to a real-time update sent from a tag on a website via a web service back to a database. All of these are legitimate, viable, and required if the resulting data is going to be a holistic view of the customer.

As you ingest data, you will inevitably be receiving it from myriad sources in many different formats (e.g., structured and unstructured), at a broad variety of cadences. This certainly sets up challenges for whoever is undertaking the process of coalescing all this data into an accessible data store. The biggest challenge is applying standard forms of data quality, as well as other forms of structuring for consistency, to the data. For example, when dealing with PII, you'll need to apply name and address standardization. That standardization must be done at the inception of the process to get the optimal benefits of that magic sequence of messages communicated to a customer at the right time and place.

All of this has to happen before you integrate the data into a data model. If it is left undone and only structured just-in-time, as some companies do, you will end up with inconsistent, imprecise results that lose your data lineage—which is increasingly important in these days of privacy regulations. And, you could potentially sub-optimize your overall targeting and messaging efforts.

Once the data is qualified, cleaned, structured, and standardized, that data—or, more specifically, the PII—is then put through a probabilistic matching process that identifies common records from different data sources. This allows the unification of these records in such a way that you get all of the information for an individual brought together from numerous sources to build the 360-degree view of the customer.

It might seem easy to just link all these records through some identification key, such as a Social Security number or some other identifier, but it's not. And that is exactly the problem a real-time approach to data matching solves. Typically, there are no common identifiers across sources in an enterprise. A record may exist in one format with one address in one system, and it could be in a completely different format with potentially

a different address in another system; they both represent the same person, but those two records look different and have no common identifier. As a result, you get the types of inconsistencies or inconsistent treatment that often occur in large organizations.

The purpose of developing a Golden Record is to bring all these records together and match them using a variety of heuristic and probabilistic methods that deduplicate the data and create a “master record.” The master record has all of the PII cleansed and formatted and structured appropriately—and all of the attributes collected from all of the contributing records associated with that master record—so that the resulting Golden Record has the individual’s identification information, including:

- **all of their various proxy identities**, whether it be a cookie or an email address or a social handle
- **all of their behavioral information**, including transaction behavior and utilization behavior
- **any KPIs, model scores, or behaviors** that are relevant to the communications you plan to have with that customer

---

The term **Golden Record** reflects how much value there is in consolidating the data in such a way that you have a complete view of a person’s interactions with a company—especially when it’s a large enterprise.

---

The quality of the interactions with individual customers improve immediately because of that knowledge, and customers can now be treated consistently with their known value to the organization and in a meaningful way that is consistent with their desires and preferences.

Think of the resulting data as an inventory of these Golden Records that represent the perfected view of the customer. But “perfected” only if the data is updated as frequently as possible. Some data elements will need to be updated on some standard batch sequence, while other data will be updated in minutes, seconds, or even milliseconds. The important thing is that the data be updated as quickly as necessary for the type of and volume of interactions that a particular brand experiences on a regular basis. When it comes to digital interactions, updating each Golden Record in milliseconds is a must.



## Step 2: Deriving Relevant Insights

The benefits of maintaining these Golden Records materialize when you use AI to derive relevant insights. But few companies have been able to scale the benefits of AI across the enterprise to the extent that you would expect given the benefits and the results it provides. This is largely due to needing highly skilled resources able to write and build models in scripting languages.

Many problems in the world of marketing, in particular, can certainly benefit from statistical modeling. Customers’ product, channel, and time of day preferences, as well as next-best product clusters and customer lifetime value, are all insights that can be important when dealing with a customer. Yet, these can get stymied because of the difficulty and cost of building those models by hand.

As such, there is an opportunity for a breakthrough that will enable a broader audience—think of them as citizen modelers—to leverage analytical and statistical techniques that will yield benefits broadly across the enterprise. The right customer engagement platform can provide this.

It is in this spirit that Redpoint Global categorizes models into four general areas (continuous outputs, categorical outputs, sequential outputs, and clustering) defined by the type of outcome one might be looking for. Outcomes typically are known as targets or signals and are the KPI or dimension of a customer that you’re trying to predict. That prediction can come in the form of a probability score, dollar value, category, or sequence types.

- **Continuous outputs:** Models that are going to result in something that looks like a dollar value. So, you might use a continuous approach to predict a customer lifetime value score or revenue score.
- **Categorical outputs:** This type of model might be a set of yes/no or dichotomous categories, or can be multiple categories; say, one of five or 10 different categories. Categorical outputs are often useful for determining insights such as channel preference or category affinity.
- **Sequential outputs:** These represent the sequencing of events or messages or content—the next-best X results. “X” stands for anything that might be in a sequence, such as content or next-best product or next-best action or the next-best anything that you will want to provide to a customer.
- **Clustering:** This is an approach to finding groupings of customers that share a similarity or similarities across a set of data dimensions that the user defines. There are numerous approaches for clustering, all of which can prove effective in finding commonalities across a customer universe that can be used to simplify the number of messages or types of content needed to effectively market.

Sequential outputs are the heart of real-time decisioning for optimized messaging. When defining a sequence of messages and optimizing those messages down to the individual level, you would use sequential tools that allow you to determine the next-best X. In this case, X would represent messages or products. If you are a product company—for example, a large clothing retailer—you might think of a sequence of products that starts with a high-priced anchor product like a pair of jeans, then additional items might sequence around it that would complement the jeans and create outfits. In a financial setting, it could be the sequencing of messages around cross-selling various products that would ensure the greatest level of loyalty and decrease the churn common in the banking industry.

Sequencing messages is fundamental to how brands build a relationship with a customer. Any given sequence might be valid for a period of time; but, based on how a customer's behavior might change, the nature or details of that sequence would then change quite significantly. So, it is important to be able to build and rebuild these models at a very high cadence to develop and optimize the sequence of products or messages for a particular customer.

The right real-time decisioning platform will allow you to develop that sequence of messages based on the circumstances, as represented in the data at the time of the model build. These models can be refreshed on a daily, hourly, or even minute-by-minute basis, depending on the velocity of interactions and the need for modification of those sequences at any given point.

It is in these analytic tools that we develop a sequence of messages that is going to ensure the greatest loyalty, greatest marketing performance, greatest level of revenue generation, and more. And it is through the sequences that you can then optimize communication with a customer, instead of going about filtering what otherwise would be a chaos-driven model of messages that may come across as random and impersonal.

Fortunately, it takes very little effort to build these models using the right AI platform and database of Golden Records. Plus, the sequences of messages or products can be developed programmatically through the system and generated for each individual customer, as a list of messages that represent the ideal for each customer.



### Step 3: Taking Action

When thinking of the conceptual cycle of data to insight to action, its application to business and marketing problems is obvious. You can use the cycle to solve business problems that benefit both businesses and their customers. In the case of real-time personalization, that action is going to be delivering that sequence of messages. This can be quite challenging for most businesses.

The idea of sending personalized messages to every single individual in a company's database would feel like an overwhelming task—one unlikely to succeed—if it's necessary to essentially build a unique set of messages for each individual in your customer base that you want to reach. This is why approaches such as clustering are so useful: You can build groupings of similar customers and use them to limit the number or variations of content you need to produce. However, if you want to move to a greater level of personalization, and resulting benefit, you need to take it down to the individual level.

#### ***So how do you do this?***

In this context, messages and channels are very tightly coupled. This is done out of necessity because executing messaging through channels is a complicated process and generalizations simplify the steps needed to make it an efficient process. But channels are really a logistical concern for the company or marketer. The message or content is ultimately what delivers value to the customer.

Real-time decisioning using the Redpoint rgOne customer engagement platform enables you to decouple the message from the channel in such a way that you can define messages that will be delivered to individual customers regardless of the channel they may appear on.

Consider an example where analysis has determined that Customer 1 should receive messages A, B, C, D—in that order. When Customer 1 visits the website, message A should be rendered once that website recognizes the customer. That seems pretty straightforward. Now imagine that an email blast is going out and Customer 1 is included in that list. The content that Customer 1 will receive is going to be determined at the time of send; therefore, it cannot be message B unless the entire audience for that email is getting message B.

What if you could change the content of that email at the moment Customer 1 opens it? Then, you could actually deliver message B—adhering to the sequence that was analytically derived. Instead of the message being preordained by channel, it was preordained by someone's identity as a customer and their preferences and delivered opportunistically wherever that customer is connecting with the brand.

Continuing the example, Customer 1 is now scheduled to receive message C. Using real-time decisioning, you can display message C when the customer opens your mobile app later that day. A visit to the website the next day prompts message D.

What you see in this example is a sequence of messages A, B, C, D delivered in a particular order based on optimizing for next-best X—and not hardwired to a particular channel. Those messages were delivered opportunistically when either a message was being sent to the customer via an outbound channel (in this case, email), or when the customer visited the brand via a synchronous interaction such as the web, mobile app, or contact center.

You can accomplish this using the rgOne platform that provides message-list intelligence along with real-time decisioning. Message-list intelligence allows you to define the sequence of messages—in this case, A, B, C, D. The platform then delivers the message using a real-time decision engine; landing that message in real time wherever a customer engages the brand. This is an enormously powerful, channel-agnostic tool for delivering exactly the sequence of messages you wish to deliver at any point in time.

This approach leaves room for variation in cases where customers take an action that might require a change in the message sequence. Because the message list is a "container" for various messages, as determined by the machine learning application, it updates programmatically as frequently as the models are refreshed. That means you can have the models refresh on any cadence or trigger event. So, if there is new data, such as new actions taken by the customer, the models are quickly refreshed, automatically, to ensure that the appropriate sequence is always at the ready, waiting for the opportunity to send the next-best message to the consumer.

Let's revisit Customer 1 and imagine that message B was a suggestion to buy a particular pair of jeans and was about to be delivered via email. But just before that message sent, Customer 1 went to the website and purchased those jeans. That behavior is incorporated into Customer 1's Golden Record in real time. The real-time decisioning engine evaluates actions every time a message is going to go out to a customer, so instead of delivering message B, you would deliver message C. This creates a better customer experience because it eliminates the frustration Customer 1 would have felt receiving a message about jeans just purchased.

Imagine never delivering a duplicate or redundant message to a customer ever again. Along with increasing customer satisfaction, you would reduce customer frustration and decrease friction in the messaging process.

## An Always-On Customer Journey

With customers' expectations for relevancy and personalization at sky-high levels, real-time decisioning should be a core element in every modern customer engagement platform today. The level of personalization that companies now need to offer across every channel has increased dramatically if they are to truly deliver the richness and flexibility that a live experience provides. Real-time decision engines are going to be at the heart of those personalized experiences. These tools will have to support a much higher level of complexity and nuance to meet the context-sensitivity of journey personalization. Supporting this takes extraordinary levels of performance because of the depth of the personalization and overall volume of customers.

Real-time decisioning allows you to deliver always-on personalization based on what's optimal for each customer, not each channel. You can automatically take a next-best action delivering the most important message a specific customer needs to see in that moment. You'll decrease generic campaigns and redundant messaging while reducing customer frustration and improving the customer journey. This simple, effective approach to personalizing customer interactions requires high-quality data, structured and maintained in real time—best managed using the Redpoint rgOne platform designed for real-time interactions.

Using real-time decisioning is an innovative way to use customer data to drive precise messaging that is automatically recalculated based on each customer's individual interactions and the overall volume of customer interactions. A true segment-of-one is now a practical reality. Ultimately, it ensures that you'll always reach the right customer at the optimal moment with the ideal message.

---

### About Redpoint Global

With Redpoint's software platform, innovative companies are transforming their customer experiences across the enterprise and driving higher revenue. Redpoint's solutions provide a remarkably unified, single point of control where all customer data is connected and every customer touchpoint intelligently orchestrated. Delivering more engaging customer experiences, highly personalized moments, relevant next-best actions and tangible ROI—this is how leading marketers lead markets. To learn more, visit [redpointglobal.com](http://redpointglobal.com).



US Headquarters: Wellesley, MA | Tel: +1 781 725 0250 EMEA Headquarters: London, UK | Tel: +44 (0)20 3948 8170  
[www.redpointglobal.com](http://www.redpointglobal.com)

© 2020 Redpoint Global Inc. All rights reserved. Redpoint, the Redpoint logo, and all Redpoint product names are trademarks of Redpoint Global Inc. All other trademarks are the property of their respective owners.