

Advanced real-time marketing with IBM Interact



How IBM Interact determines the right message to present in inbound marketing channels

Contents

- 1 Presenting personalized messages in real time
 - 2 Interact's real-time decision process: A high-level overview
 - 4 A deeper dive: How Interact works
 - 7 Self-learning arbitration
 - 10 Other advanced features
 - 12 An architecture design for cross-channel marketing
 - 16 Architected for high performance
 - 16 Distinguishing features and benefits of IBM Interact
 - 18 Conclusion
-

Presenting personalized messages in real time

Customers have grown increasingly overwhelmed with outbound marketing offers, and, consequently, the effectiveness of traditional outbound marketing campaigns has significantly declined. Marketers are now recognizing the advantages of connecting with customers on their own terms—during those moments when customers and prospects contact them.

But how can marketers apply the same analytic rigor to marketing through inbound channels as they typically do for outbound campaigns? How can inbound marketing and personalization factor in both what's already known about the customer and what's just been learned during the most recent interaction? And how can anonymous web visitors be marketed to in real time in the same way as known customers?

IBM® Interact allows marketers to personalize in real time the experience of customers interacting with websites, call centers and other inbound marketing channels. Interact uses powerful, behavioral targeting analytics and marketer-defined business logic to deliver the optimal marketing message in each channel. Combined with IBM Campaign, Interact enables inbound and outbound marketing to be coordinated into a single strategy for building customer relationships, using a single technology platform.

Take full advantage of inbound customer contact

IBM Interact offers a variety of marketing management features:

- **Presents targeted personalized marketing messages in real time through interactive, inbound channels:** Personalization is based on behavioral targeting for anonymous visitors with improved targeting for known or returning customers that will combine historical customer data and current behavior to derive the best possible communication. For known customers, previous marketing events will play a role in how they are treated. What we learn in each customer interaction informs future communication for an individual customer and across the customer base.



- **Plugs into any customer touchpoint where inbound contact or interactions may occur:** Interact works with any inbound touchpoint — website, call center, kiosk, point-of-sale, mobile app., etc.
- **Highly manageable, real-time decisioning logic:** A combination of segmentation and business rules enables marketers to easily define strategies for interacting with customers and prospects and, just as importantly, easily modify those strategies and understand the impact of changes.
- **Self-learning and arbitration:** A configurable Bayesian algorithm combined with integrated offer arbitration capabilities builds insight over time and helps determine the best offer or message to present to each customer or prospect. Self-learning extends and integrates with any existing predictive modeling capabilities used by the organization.
- **Role-based user interfaces (UI):** Separate user interfaces for marketers, technical analysts, and operational managers of interactive channels create an uncomplicated and intuitive user experience for each person involved in the inbound marketing process.
- **Insightful reports and analyses:** Interact provides marketing performance analysis, interactive channel activity analysis, campaign strategy inventories, and deployment auditing—all intended to improve managerial oversight and marketing results.
- **Integration with outbound marketing capabilities:** Interact is part of a single cross-channel marketing solution, at the heart of which is IBM Campaign, enabling inbound and outbound marketing to be coordinated into a single strategy for building customer relationships, using just one technology platform.
- **Performance optimized run-time engine:** To maximize performance, Interact is Java-based, features configurable multi-threaded, horizontal-scalable processing supporting redundancy, high-volume throughput and extremely low latency for the derived marketing decisions.
- **Meets the most rigorous performance requirements:** Interact has been architected for high performance even when facing the heaviest volume of customer interactions, with the ability to manage thousands of transactions per second with sub 200 millisecond average response time.

With Interact, marketers can address the key challenges facing today's organizations: how to take advantage of inbound marketing opportunities and how to identify what the next best offer or message or communication is for each customer. The following discussion will provide a deeper look into what Interact is and how it enables real-time marketing.

Interact's real-time decision process: A high-level overview

The starting point of Interact's decision process occurs when the inbound touchpoint (be it a website, the call center application, mobile app, ATM, etc.) makes a call to Interact, indicating that a real-time decision is required, for example, to determine an offer to present in an empty space on a Web page. As part of this call, the website passes the context of the interaction with the customer or prospect to Interact—information about what the customer has been doing during that interaction (e.g., what website pages the customer has visited). At the same time, Interact is able to access profile data about the customer or prospect—often the kind of information that sits in a customer database. Think of the profile data as what was known about the customer or prospect before the interaction started, and the context as what has been learned about the customer in the last few moments.

With this information, Interact makes a real-time decision about how to personalize that interaction—for example, what offer or offers to present. In making this decision, Interact uses a combination of segmentation logic and rules—which would have been manually defined by the marketing team—and algorithms, which are set to monitor interactions over time and automate the process of building a mathematical model that can help make personalization decisions.

As a result of that real-time decision, Interact passes information back to the website, most likely the right offer or offers that have been selected by Interact to be presented during that Interaction.

Having described how Interact works at high level, let's see an example of how it works in practice.

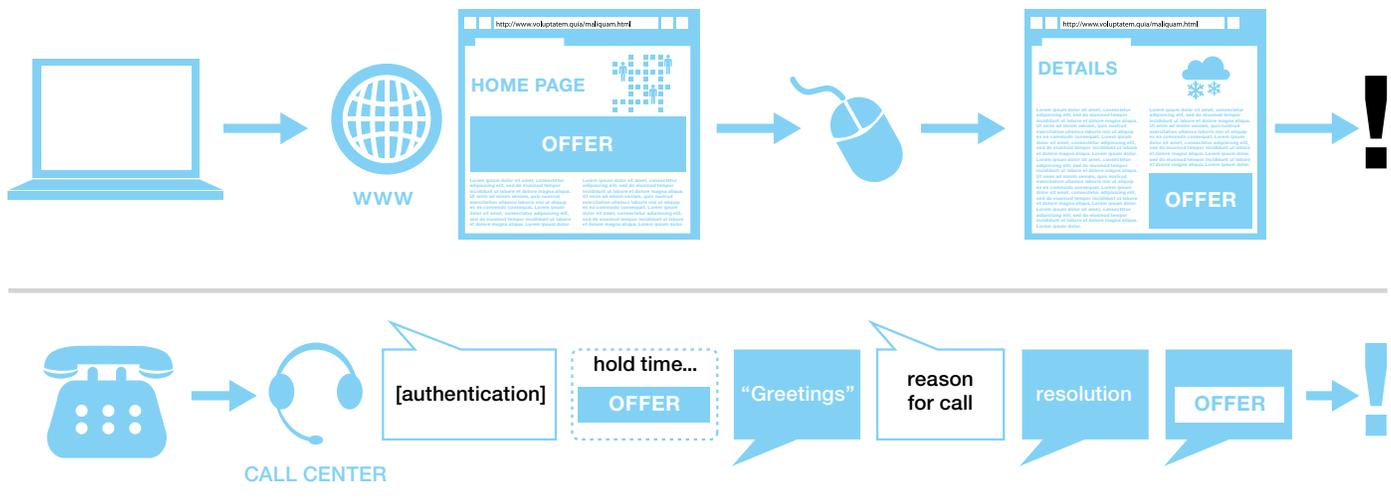


Figure 1: Interact enables “inbound marketing”

Imagine a customer accesses your website via the Internet (see figure 1). The moment the person hits your home page is your first opportunity to conduct inbound marketing. Perhaps there’s a spot on that page where a personalized offer could be presented. Interact can decide what offer or message should be presented there. Perhaps the website visitor accepts the offer. Or that person may not accept the offer but navigate around on your website for awhile. Every time the customer clicks, Interact is learning more about the person’s interests, needs and desires. Deeper on your site, on pages where there may be more details about various topics, there may be other spots where personalized offers or messages could be presented. By combining what you already know about the visitor (customer profile data) and what you’ve just learned from their browsing behavior (behavior attributes), you can reconsider, in real time, the offer you want to present to them. Perhaps this time the visitor is more interested in the offer chosen for them because it’s more relevant, and accepts the offer—the result is a successful conversion.

The same scenario can occur in other channels. If a customer picks up the phone and calls your call center (figure 1) they may first authenticate or identify themselves using the automated attendant, and then sit on hold for a few moments waiting for a call center agent to become available. During this hold time is the first opportunity for inbound marketing, and Interact could select from a set of offers to present (using a recorded voice) to the caller while on hold. Later, when the call center agent picks up, the agent will talk to the caller and learn why they’re calling, hopefully addressing the caller’s need or resolving their issue—and perhaps at the same time learning valuable new information about the customer. If Interact is connected to the call center application the agent is using, it can take that just-learned information about the customer and immediately recommend a new offer to present to the customer verbally. This new offer may be more relevant because of the new information learned during the call, and thereby achieve another successful conversion.

In the end, what “inbound marketing” really means is turning every customer touchpoint into a channel for the kind of personalized marketing messages previously only delivered through outbound campaigns.

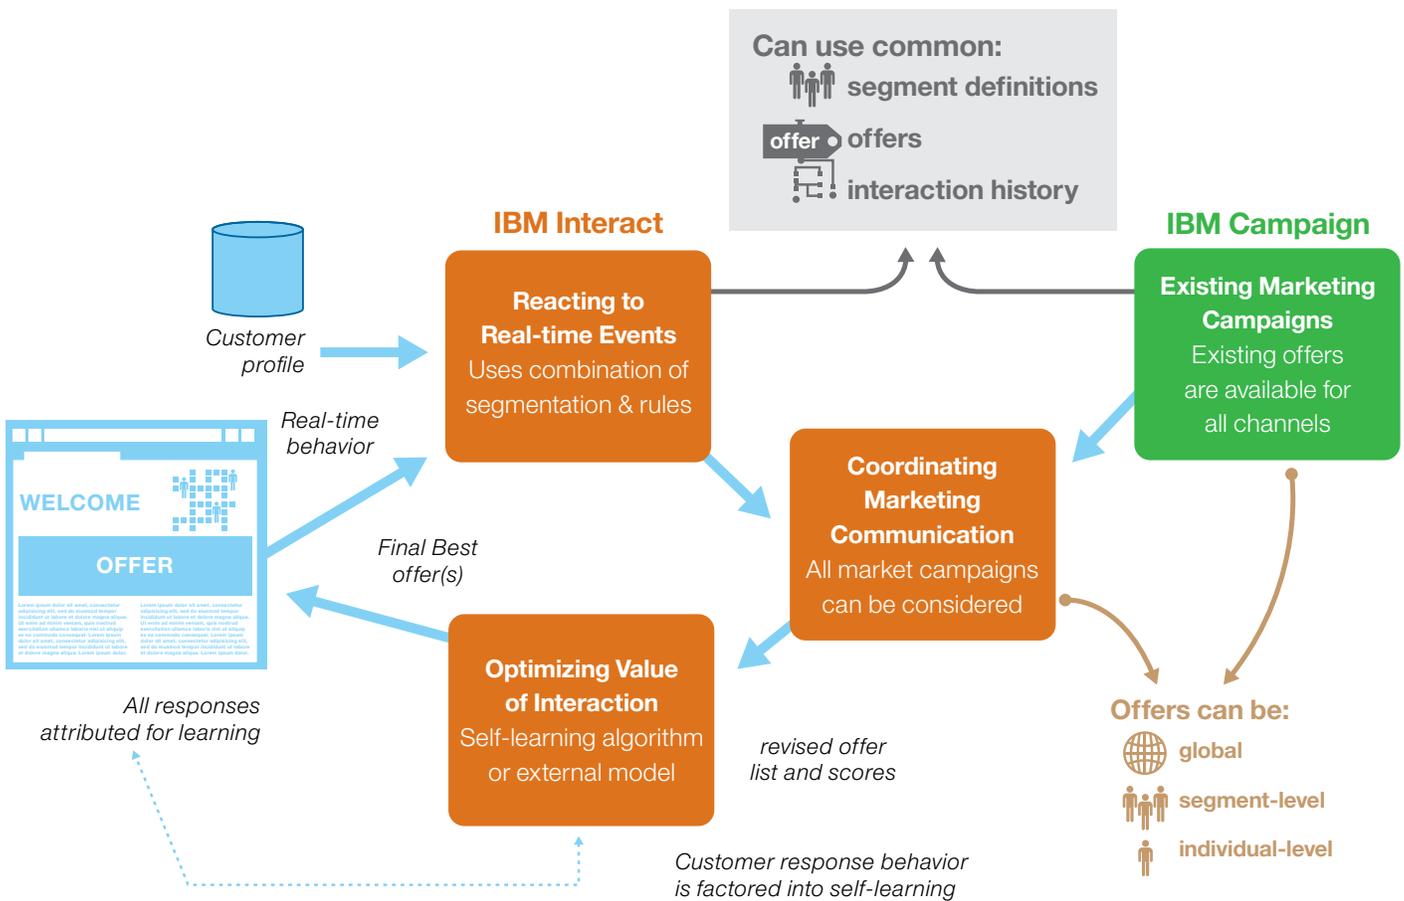


Figure 2: How IBM Interact works — detailed

A deeper dive: How Interact works

Reacting to real-time events

As mentioned in the previous example, the starting point of Interact’s decision process is when the inbound channel—the website, the call center application, mobile app, Facebook page, ATM, or whatever the customer touchpoint is—makes a call to Interact. Through Interact’s API and various integration methods, Interact receives the “real-time context.” By “real-time context” IBM means information about what is happening during the Web visit, information which Interact will use to make a decision about what personalized offer to deliver. It’s important to note that this passing of real-time context is not a one-time event. Throughout a Web visit (or a call center call, ATM or mobile session, etc.), the website (or other system) can pass new information to Interact continuously, ensuring Interact has a complete picture of what has happened during the visit up to that very instant.

The real-time context is only one part of the basis on which Interact makes its real-time decision. Interact can also use a customer profile, optimized for usage in real-time scenarios, which contains customer data considered relevant to real-time decisions. Think of this as what you knew about the customer before the current session. It may contain demographical information, segment categorization, product ownership data, purchase history, etc. This information typically sits in a local data store that is created just for Interact to use, although Interact can also retrieve this information from a web service call (e.g. master data management system). Interact can make any number of real-time call-outs to any Web services that can deliver other relevant data instantaneously. For example, Interact might need to know someone’s up-to-the-minute credit score.

With all of this data made available, the first part of Interact's real-time logic can execute. This logic uses a combination of a traditional segmentation approach and rules to identify candidate offers for the person in question and score those offers. Segment definitions span rich customer data, behavioral events in the touchpoint and patterns of events that occur across touchpoints. This blend of segmentation and rules provides marketers with a lot of control over the real-time logic, and makes it easy for them to set up the logic using concepts familiar to them. The output of this part of the process is a candidate list of scored offers that are eligible for the specific customer and real-time situation. But creating this list of candidates is only the beginning of the story of how Interact makes decisions.

Coordinated cross-channel marketing

Interact is part of IBM's Enterprise Marketing Management suite, and so its real-time decisions can be easily connected and coordinated with what's happening elsewhere in your overall marketing efforts. One example of this integration is that Interact can accept additional scored offers that you have pre-calculated in IBM Campaign. That is, you may have decided before the current live interaction what offers you want to be sure are always considered for a given person during a live interaction—we call this a "white list"—or quite the opposite, you may want to make sure the person never sees certain offers—we call this a "black list."

This kind of pre-calculation often takes place when planning and executing outbound campaigns, so this is an illustration of how IBM's inbound marketing capabilities are connected to IBM's outbound marketing capabilities. But inbound/real-time and outbound/pre-calculated aren't just connected by the fact that Interact can work with both types of offers. Both decision processes—real-time and pre-calculated—can leverage a common set of building blocks, such as segment definitions, offers and an interaction history (interaction history being a history of contacts and responses and the ability to attribute accurately the former to the latter). This is a part of what enables IBM's Enterprise Marketing Management suite to facilitate truly cross-channel customer dialogues, because every marketing activity, no matter what channel it happens in, be it an inbound or outbound channel, is managed using the same system.

Optimizing the value derived from the customer interaction

Once Interact has accessed the real-time offers and the pre-calculated white lists and black lists, Interact undertakes a process of removing offers from its lists of candidates through a set of internal suppression rules. An example of its own suppression rules might be that Interact knows there's a limit to how many times a particular offer should be presented to a person, and so if a particular offer has been shown too many times already, Interact will remove it from consideration even though it may have been put on the initial list of candidates by either the real-time logic or the pre-calculated white list. The output of this adding & removing process is a revised list of scored offers.

There is tremendous flexibility in the nature of this revised list of scores offers. In particular, there are many options when it comes to the audience level at which they are targeted. For example, offers can be put on the list because they are global offers—i.e., everyone is eligible for them. Or they can be segment-level offers—i.e., an offer is added to the list of candidates for that person because the person has been identified as being in a particular static or behavioral segment. And, of course, there can be individual-level offers, where an offer has been selected as a candidate for a specific person based on their particular characteristics. This interplay between global, segment-level and individual-level offer eligibility is part of Interact's real-time logic working in conjunction with pre-calculated decisions to build the ultimate list of offers to be consider for a single individual.

The final major step in Interact's process is the arbitration step. During arbitration, Interact takes the final version of the offer list with scores, and uses a self-learning algorithm to adjust those scores. Note that self learning must "train" up to a user-configured confidence level before it is allowed to substitute its judgment over the marketer's. Interact's self-learning algorithm watches, over time, all the offers presented and the response or non-response, and uses the characteristics of responders to predict other people's response likelihood in the future. Interact comes with its own self-learning algorithm, or an outside model can be used. It is during this arbitration step that the science of this predictive self-learning model is blended with the art of the logic created by the marketer and used during earlier steps in the process to choose the optimal offer or offers from the list of candidates.

It is important to mention that marketers have the option to “turn off” self-learning if they wish. Or they can choose to run self-learning in a mode that never changes offer scores, but rather simply allows the marketer to view reports that demonstrate how much more effective an offer would be if self-learning was active.

Once arbitration is complete, the final offer(s)—depending on how many offers have been called for—are passed back to the customer touchpoint.

It’s important to remember that everything just described takes place in a fraction of second—the specific speed being determined by the complexity required for the decision and the hardware resources deployed to support the Interact runtime application.

Machine-based learning in real-time

The story is almost over at this point, but not quite. Interact will record the fact that an offer has been presented, and also the response or rejections of offers served. This information not only gets factored into the self-learning algorithm for future scoring, but also can be factored into future batch campaigns by becoming part of the centralized interaction history that is used by both Interact and Campaign.

In figure 3, you can see a high-level illustration of how Interact might work on a Web page. In the middle of the page you can see that there are two “zones” on this Web page that need a personalized message. The Web page calls Interact and Interact runs a flow chart, which is part of the real-time logic a marketer has defined to determine what to present to the

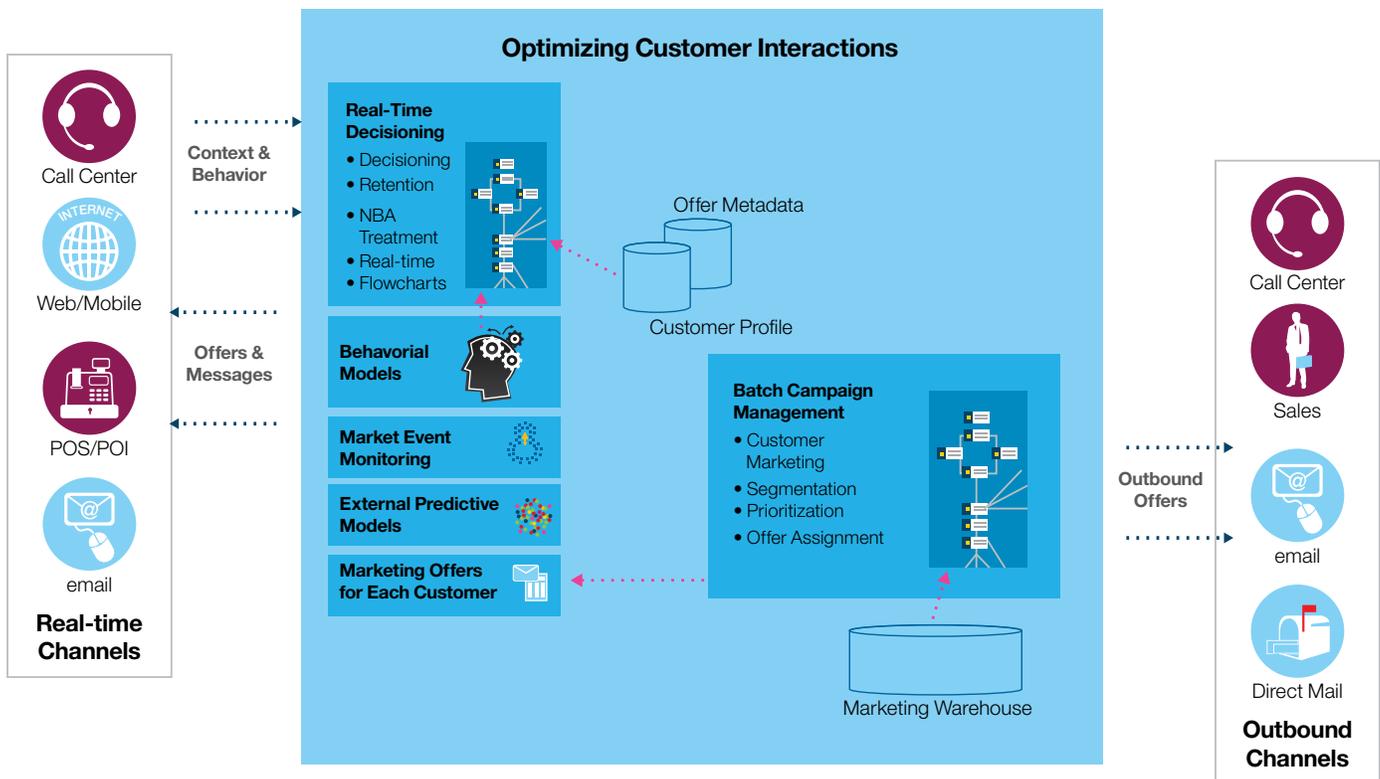


Figure 3: Optimizing each customer interaction

Web visitor. The flow chart will use the marketer's logic to classify the customer (based on historical data) and the customer's behavior (based on the customer online actions). The individual customer will be classified into segments based on history and behavior. In our example, the Web visitor winds up in three different segments, based on his or her behavior on the site and other information known about the visitor. Based on that segmentation, Interact identifies candidate offers for each of these segments and calculates a score for each offer for that specific individual. Note that in this case the "offer" may not be a promotion, rather just an identification of the best content for this person. After black and white list consultation and arbitration, the best offers for this person are presented on the website. Beyond that, Interact could adjust what is presented on this page based on further behavioral cues from the person's browser behavior. And Interact could be configured to do A/B testing to test various offers over time.

Self-learning arbitration

Most learning systems use a formula known as the Bayesian model which is designed to optimize the acceptance rate. There can however, be problems with this approach if used solely on its own. Straight Bayesian models will learn very quickly with very limited data which can be desirable, but also can cause overwhelming bias. This occurs when quick learning over-estimates the probability of a customer response if the first few customers react in a certain way.

For example, imagine you were visiting an online news portal and clicked on an article about a bank robbery—not because you are interested in crime, but because the article quoted your friend who was interviewed for the story. Now imagine that for the next three weeks you were targeted with stories about crime, a topic in which you have no interest in at all. The model has actually made a false assumption and, consequentially, targeted you with unattractive offers.

To prevent overwhelming bias the Interact learning models perform a confidence level test that tests the statistical relevance of every model variable before that variable is used to predict the outcome. This is tunable within the algorithm so you can adjust how responsive the models will be, while guarding against overwhelming bias.

The other potential issue with the Bayesian model is the idea of over learning. For example, some trends may be very hot but very brief, such as the launch of a new iPhone or a new music album. The Bayesian model alone doesn't account for this concept of short-term spikes in behavior. It is imperative that as response rates shift a learning model also shifts. To combat this Interact uses configurable time intervals that will automatically increase the weight of more recent responses.

Lastly, most predictive modeling approaches calculate the likelihood of response, not the value of the ultimate response. Interact uses optimized offer prioritization on top of predictable and management segmentation strategies to optimize the combination of the value of the offer with the likely acceptance rate.

For example, consider a bank stands to benefit tremendously if a customer refinances his or her mortgage. While the likelihood of the customer accepting the offer may only be 10%, the bank stands to make over \$10,000 in the next three years. While that customer may be more inclined, say by 25%, to respond to an offer for a new credit card, the bank only stands to make \$500 over the next few years. So even though the likelihood of that customer accepting the refinance offer is lower than the new credit card offer, the potential value of the refinance offer is far greater for the bank.

Interact is unique in that it takes a combination of the likelihood that someone would accept an offer and the value that would come with that accepted offer and uses that for arbitration. Those factors are boiled down to scores which are then multiplied to get a net present value of the offer that is most optimal for a given individual.

How self-learning arbitration works

The following example demonstrates Interact’s self-learning arbitration model in action. In figure 4, you can see that Sally Jones has been served a 3.9% balance transfer offer via a Web page. She accepts the offer and Interact begins to collect statistics about her persona based on a series of attributes of interest. In this case, she makes over \$100K a year, she has a \$5K balance, she owns a money market account, and she is a female. Over time, we will collect the same information for all visitors—also making note of how many people accept versus those who have ignored or rejected the offer.

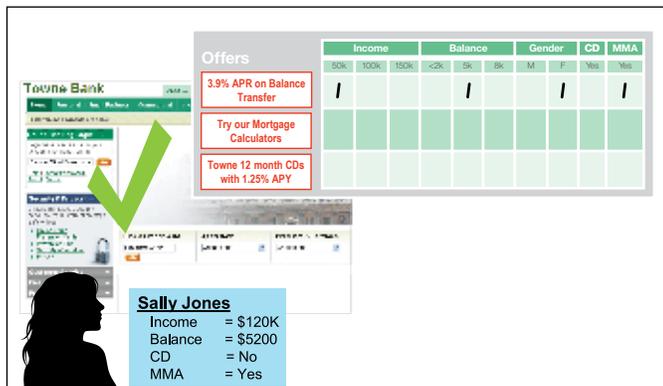


Figure 4: Gathering statistics

In figure 5, we see that John Smith is also online. He has been served the same offer as Sally, but he ignores it.

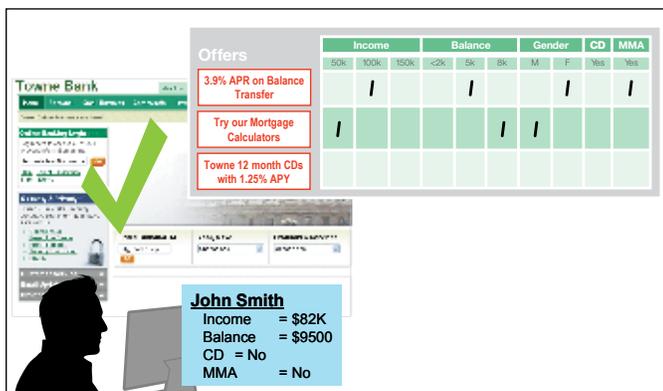


Figure 5: Each offer is treated independently

The system makes note of that and after time passes it serves a new offer for a mortgage calculator which he does accept. Interact records his statistics—he is a man, with an income greater than \$50K, he has an \$8K balance and he owns a money market account.

It is important to note that offers can span channels and channels can share learnings. A call center representative can offer the same balance transfer or mortgage calculator offers to a caller that Sally and John responded to.



Figure 6: Patterns emerge based on the masses

Over time, as more and more people respond to the offers, specific attributes emerge as predictive. Patterns begin to materialize and we can easily identify the attributes of those visitors most likely to respond to an offer. For example (figure 6), with the 3.9% APR on Balance Transfer offer we see that the people that accepted typically had incomes around \$50K with an \$8K balance—those are the predictive attributes. Something else we have learned is that some of the attributes don’t tend to make a difference, for example if a visitor is male or female. There is no entropy in those variables so we can choose not to use gender as a learning attribute, at least for this offer. Also, if there is not yet enough data on a specific attribute, for example the Money Market Account (MMA) attribute for the 3.9% offer, we won’t use that variable yet because we don’t want to create bias. However we will continue to collect data on that attribute.

Learning models will store a number of attributes for each offer the system will learn on, typically anywhere from 30-60. These are typical attributes associated with the customer or specific behavior the customer is exhibiting at the time the offer is presented. You do not have to use the same attributes for all offers. Rather, each offer or sets of offers can be customized to learn on a unique set of attributes that are relevant to that specific offer set. And marketers have the ability to add new attributes anytime during a campaign. Interact will immediately begin collecting data on those new attributes, but the model won't exploit the information until there is statistical relevance for that attribute. Once there is, Interact will begin factoring in those attributes as part of the self-learning process.

It is important to note that unlike many learning products, Interact's self-learning capability factors in both customer attributes (i.e. male/female, income level, etc.) and behavioral attributes (web browsing history, time of day, etc.), providing marketers with a more comprehensive and holistic view of how visitors are responding to offers. Also, self-learning can be applied after a granular segment level rather than global

segment level. For example, an offer may be very successful for people who drive red Ferraris and have blue eyes, but a total failure for people who drive Fords and have green eyes.

Reporting

One of the complaints customers often make of learning systems is that they are a mysterious black box, offering no visibility into what is really happening. Interact comes with a set of reports that provide you direct visibility into what affects learning is actually having on offers.

One example of a highly useful Interact report is the Self-Learning Model Details report which allows marketers to visualize what attributes were most predictive for a given offer. In the case of figure 7, we can see that the people most likely to respond to this offer had a significant savings balance, bought CDs in the past and had visited the CD Web page. The report identifies the pattern of what's important for each offer and exposes that data to the marketer. This is valuable information that they can then take to inform their own offline modeling in the future—for both inbound and outbound campaigns.

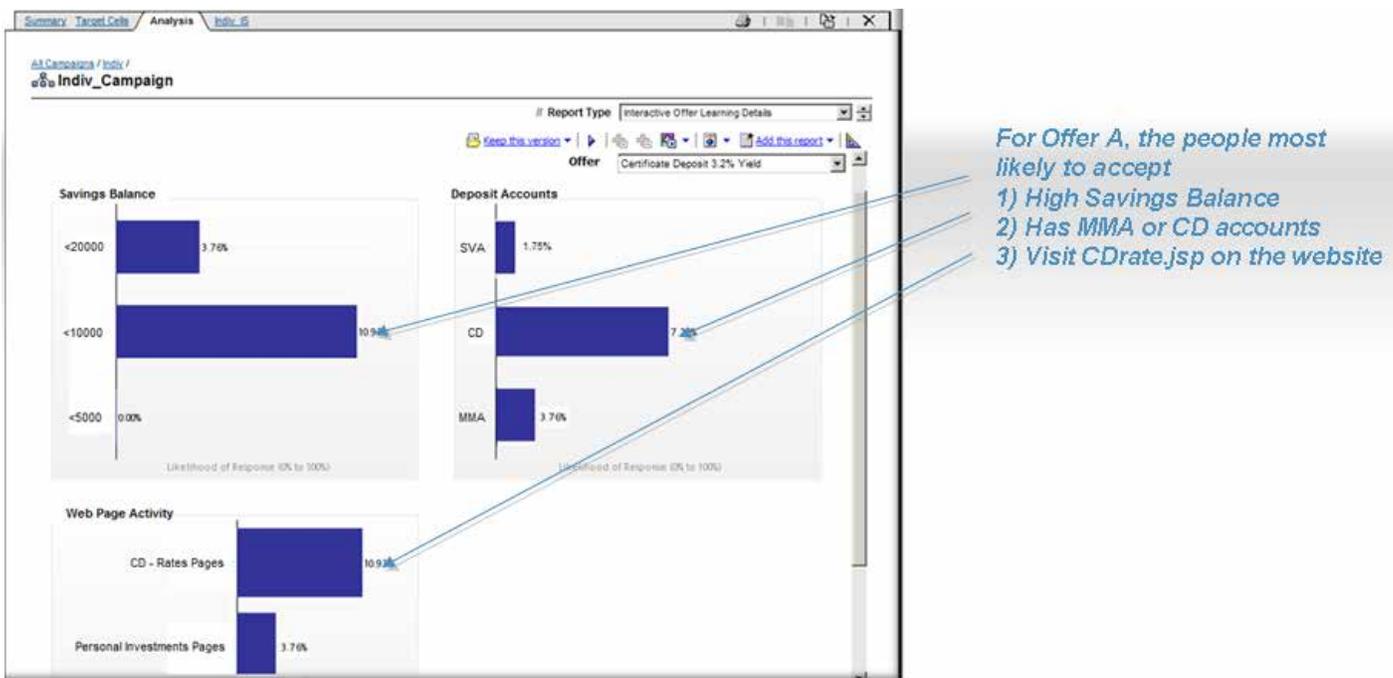


Figure 7: Self-Learning Model Details Report



Figure 8 Self-learning Lift Report

Built into the system, Interact offers marketers the ability to serve offers without learning as control so they can measure how well the offers are performing both with and without learning. Figure 8 is a Self-learning Lift Report that shows the acceptance rate of offers, comparing offers presented by random and offers presented by Interact built-in learning. This report shows these values over time, so you can see trends for the better, or worse, and use that information to refine your marketing strategies. In this case, the blue denotes the cumulative random acceptance rates and the green denotes the cumulative model-driven acceptance rates. The marketer can clearly see in this instance that serving offers using learning vs. a random approach had a significant positive impact on acceptance rates.

Other advanced features

Geo targeting: Interact enables inbound real-time or event-triggered outbound tactics to allow marketers to prioritize offers close to the customer’s current location. For example, marketers can customize the system to only serve offers within a specific radius—which is key for a mobile customer. For example, using the same eligibility rules and scoring rules a marketer may have two competing offers, one for Fashion Retailer 1 and one for Fashion Retailer 2. Even though the customer may prefer Fashion Retailer 2 based on previous behavior, the geo targeting functionality of Interact knows that Fashion Retailer 2 is five miles away and Fashion Retailer 1 is across the street, so the likelihood is that the customer will be more likely to accept the Fashion Retailer 1 offer.

Offer marketplaces and Offers By SQL: Particularly useful when serving thousands of potential offers, Interact enables categorical selection offers. Rather than setting up rules (if x type of person visits or performs x kind of activity), marketers can put offers into an offer list and leverage queries. So when a customer comes into a Web page, for example about TVs, Interact will perform a query of all of the potential offers relevant with TV, and choose those that score highest for that customer to present. This pre-filtering process means that scoring, constraints, learning, etc. only happens on relevant offers. So instead of dealing with thousands of potential offers, the pre-filtering process can narrow that down to tens or hundreds, improving reaction time to within fractions of a second.

This feature is called Offers by SQL because marketers can actually input the query logic that will pull the offers out of that list, making it highly flexible. Instead of a product category, marketers can also set up categories based on a customer segment (i.e. high affluence, regional, etc.), making the system highly flexible regarding which offers you can serve.

Offer learning granularity: Another flexible feature of Interact is that it enables granular learning by allowing different sets of offers to learn on different sets of attributes. This is particularly useful for companies with very diverse customer bases. Take Honda, for example. The global company is best known for selling cars, but it also has a marine division. Learning on a vehicle type would not make sense for those marketers responsible for promoting outboard engines.

Offer service limits and constraints: Interact gives marketers additional control by allowing them to build in limits and constraints for an offer, for example: “do not serve this offer more than x times an hour, in total, in one month, etc.”

Version control and roll-back: Marketers can be confident in their ability to manage rules or contact strategies, knowing that if changes were made and conversions dip instead of rise, they can easily roll back to the more successful version they had the day before.

Offer parameterization: Interact allows marketers to vary the parameters on an offer to optimize their response and minimize their risk. For example, imagine you are planning to give 10 customers an offer with a 0% rate for balance transfers for 6 months that will then switch over to 9.6%. As a marketer, you may choose to deliver different rate offers based on the risk that customer poses. If someone is subprime your risk is higher and you will want to offer a higher interest rate vs. someone who poses less risk. So while the offer stays the same, the marketer is able to modify the parameters to optimize the outcome.

Customer profile services: Interact has the ability to integrate with any MDM server or SOA architecture to pull customer data used for segmentation and eligibility. While the traditional approach is to quickly pull the customer details from an operational database, Interact can use your SOA service to access this information.

Event pattern detection: Most systems can react to a single event and have some rule that changes the content or marketing message served. One of the challenges marketers have is measuring a customer’s interest or behavior over time. Interact adds an engine that will look for patterns of behavior that may indicate interest. For example, say an investment customer looks at two 529 college savings pages and an educational IRA page. This would be an indication of strong interest in a tax deferred educational saving for their children or grandchildren. In this case the pattern would tell Interact to increase the value and score of offers and messages associated with these types of products. Maybe this strong interest would drive decisions toward requesting a direct customer communication with an educational saving specialist. Key to this feature is that these patterns can span real-time sessions and span channels.

An architecture designed for cross-channel marketing

Interact consists of two major components, the design environment and the runtime environment. And some customers may have optional testing runtime environments as well. The following figure (figure 9) shows the high-level architecture overview.

The design environment is where you perform the majority of your Interact configuration. The design environment is installed with IBM Campaign and references the Campaign system tables and your customer databases. IBM Campaign and IBM Interact share key marketing components, including offers and cross-channel Interaction history.

After you design and configure how you want Interact to handle customer interactions, you deploy that data to either a

testing runtime environment for testing or a production runtime environment for real-time customer interaction.

In the design environment, you define what Interact does at certain points in your touchpoint by configuring interactive channels. You then divide your customers into segments by creating interactive flowcharts. Within interactive flowcharts, you can perform test runs to confirm that your customer data is segmented correctly. Next, you define offers and then assign the offers to segments within an interaction strategy. Once you have configured all of the Interact components, you are ready to deploy the configuration to a staging runtime environment. While offers are assigned to individuals based on customer behavioral segments definitions, the actual end offer that is served is based on a comparative score of the offer that is specific to that individual (see advance features and discussion on learning above.)

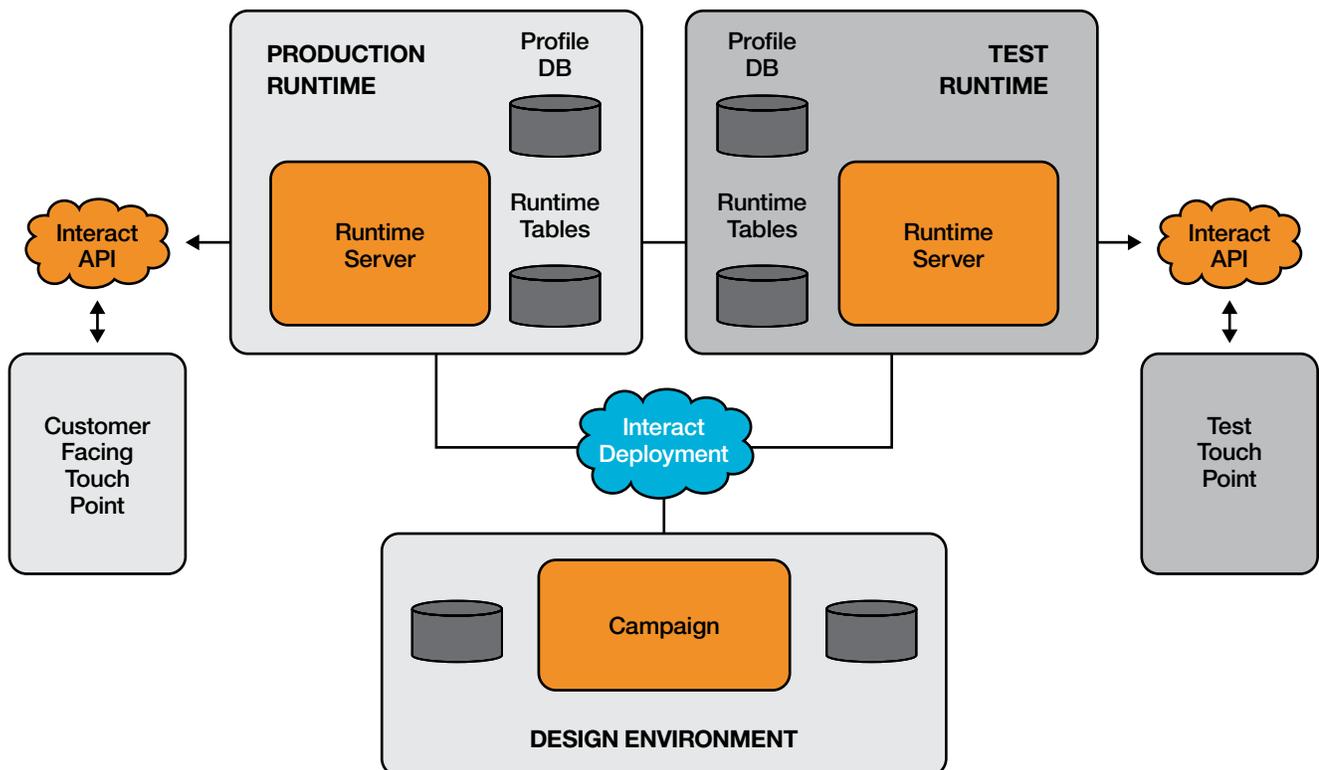


Figure 9: IBM Interact architecture — high-level view

Interact deployments consist of the following:

- Interact configuration data including interactive channels and interaction strategies
- A subset of IBM Campaign data including smart segments, offers, and interactive flowcharts

In the staging runtime environment—which is the same as a production runtime environment except that it is not customer-facing—you can test the entirety of your Interact configuration, including the API integration with your touchpoint. During runtime, a customer—or in the case of a staging server, someone testing the system—takes actions in the touchpoint. These actions send events or requests for data to the runtime server by means of the Interact API. The runtime server then responds with results, such as presenting a set of offers (data) or re-segmenting a customer into a new segment (event). You can continue modifying your Interact configuration in Campaign and re-deploying it to the runtime environment until you are satisfied with the behavior. You can then deploy the configuration to the production runtime environment.

The production runtime servers record statistical and historical data such as contact history and response history. An Interact design time utility copies the contact history and response history data from staging tables in the production runtime server group to your Campaign contact and response history. This interaction history is common across Campaign and Interact. This data is used in reports that you can use to determine the effectiveness of your Interact installation and revise your configurations as necessary. This data can also be used by Campaign and other IBM products such as IBM Contact Optimization, integrating your real-time campaigns with your traditional campaigns. For example, if a customer has accepted an offer on your website, you can use that data in Campaign to ensure either that the same offer is not sent by mail, or that you follow up the offer with a telephone call.

Horizontal Scalability

As shown in the diagram below, Interact has been design from the ground up to support multiple run-time servers and

multiple server clusters. This helps support high-volume throughput, low latency for response and local and global redundancy for mission-critical systems.

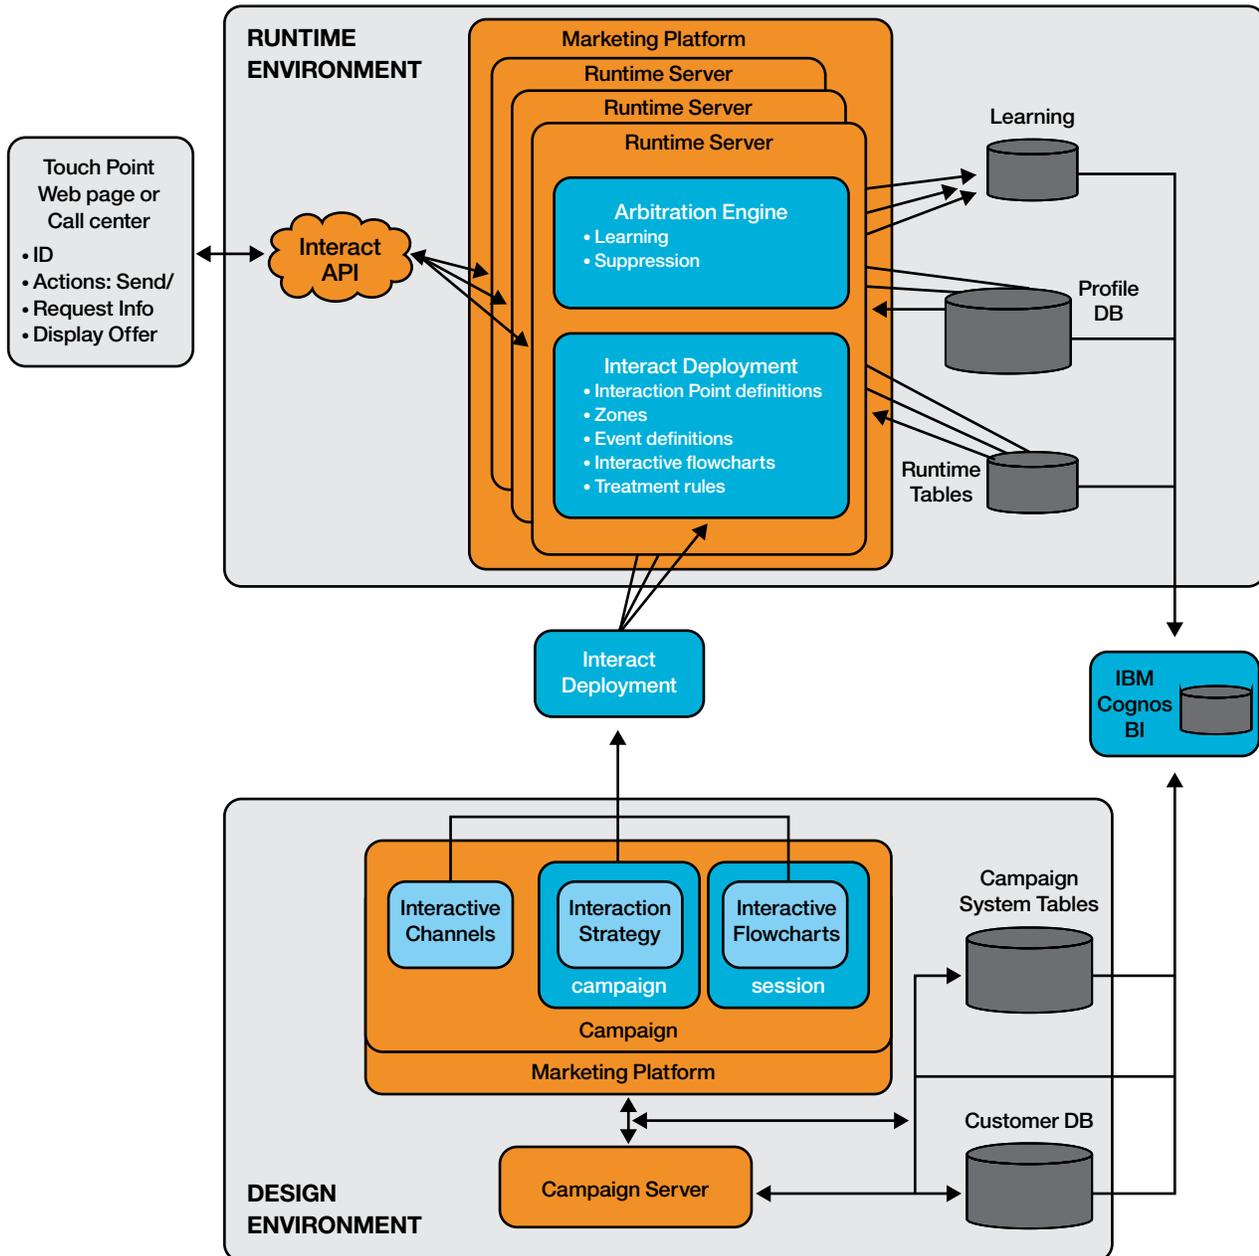


Figure 10: Interact architecture – in detail

Integration with IBM Campaign

Interact was built as part of a larger cross channel marketing strategy. Used in conjunction with IBM Campaign, marketers can coordinate their outbound communications with inbound decisions, and also influence future outbound communications based on inbound interactions. Interact and Campaign share key building blocks such as segments, offers, interaction history, attribution and even optimization. Working together, these two products create a true cross channel-experience.

One of the benefits of Interact is that it does not require a large, dense data warehouse for the segmentation and arbitration process. Rather, leveraging Campaign or another MDM, Interact is able to use a thinner layer of the customer

profile and apply behavioral learning in the decisioning process. Interact can combine pre-calculated offers created in Campaign, with offers it calculates in real time, to determine the best offer, topic or message to push to the customer.

It is important to note that in addition to producing outbound campaigns—like email and direct mail—IBM Campaign can stage offers as well. Interact can include those scores and bring those offers into Interact’s arbitration service to select the best possible message or offer. Interact can also write information back to Campaign—for example, if a customer requests more information during a session monitored by Interact. Campaign can then leverage that information for future outbound marketing initiatives.

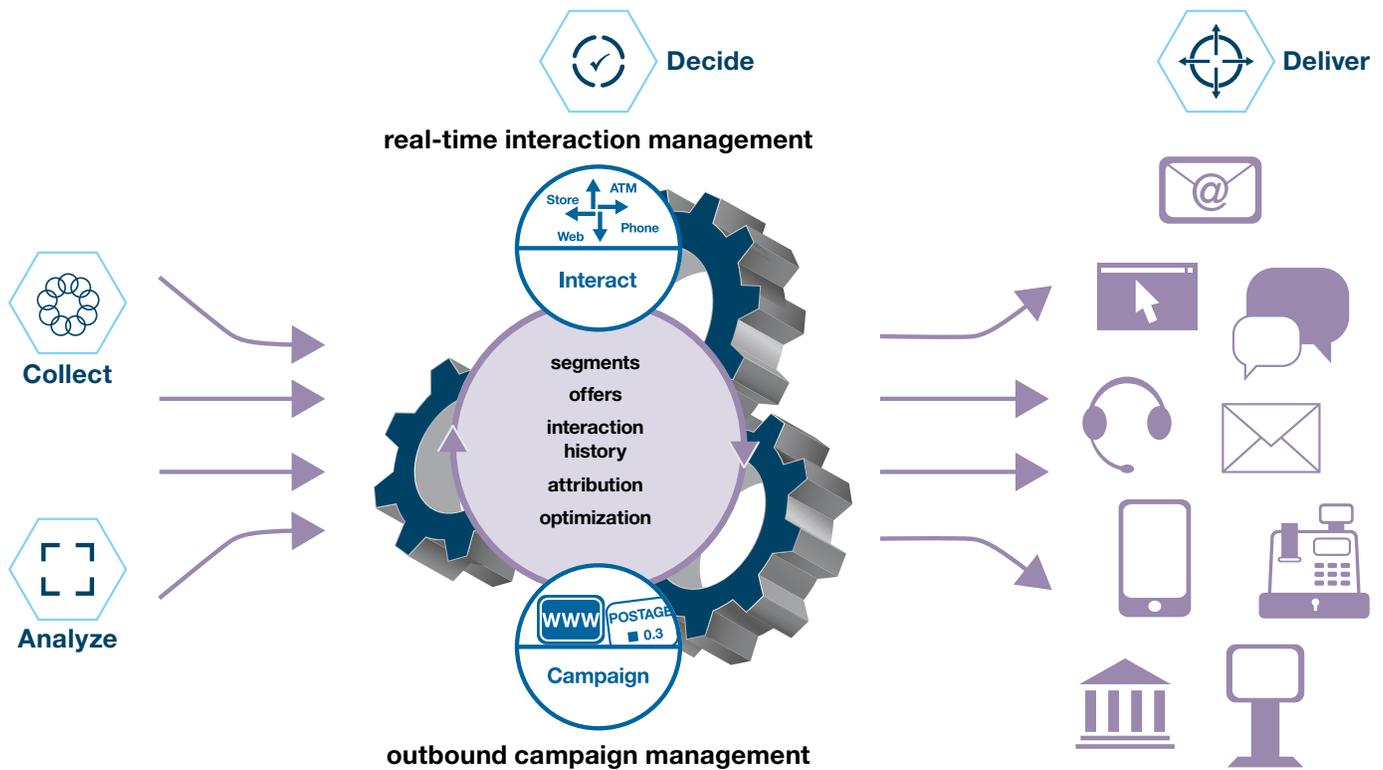


Figure 11: Together, IBM Interact (inbound) and IBM Campaign (outbound), offer integrated, centralized decision-making capability

Architected for high performance

IBM Interact has been architected for high performance, even when facing the heaviest volume of customer interactions. Leveraging its Java runtime engine and architecture for horizontal scaling, IBM Interact can consistently deliver thousands of transactions per second with response times within milliseconds and when configured optimally for each deployment.

Example Interact benchmark performance metrics include the ability to scale to meet the most rigorous performance requirements:

- 100,000+ of concurrent sessions
- 500-1000 transactions per second and up
- Overall average latency between 50 and 200 milliseconds
- Horizontal scaling tested across dozens of servers in a cluster

Distinguishing features and benefits of IBM Interact

- **Designed with the marketer in mind:** Unlike other learning solutions that are based on generic rule engines and require a lot of IT intervention, Interact was designed for today's marketer, with built-in marketing processes and concepts—like offer, suppression, customer segments (historical and behavioral), sampling/testing, etc. And Interact gives marketers more control over the arbitration process through a full set of UIs based on constructs marketers are familiar with so they can manage the system themselves.
- **Interact works across multiple channels:** Interact works with any inbound marketing touchpoint (web, mobile, Facebook app, call center, IVR/VRU, ATM, kiosk, etc). It is the cross-channel, single source of truth for next best offer and the only real way you can coordinate and optimize across all channels.
- **Interact provides far greater personalization:** Whereas most learning systems are solely rules based, generating decisions based on the “wisdom of the crowd”, Interact offers a far more personalized approach to inbound marketing. With the ability to factor in both personal and behavioral attributes of the individual in combination with targeted based customer data, pre-calculated and real-time offers, Interact's inbound offers can be uniquely tailored to the individual.
- **Interact is part of a full marketing suite ecosystem:** Interact lives and breathes in an ecosystem of a marketing suite, and can therefore take advantage of the entire system of engagement (leads, campaign, event detection, etc.) versus other solutions that focus on just one part of the marketing problem or a single channel.
- **Interact is the most mature, out-of-the box solution on the market:** As you analyze the numerous advanced features Interact has to offer, you'll recognize that tremendous thought, time and resources have gone in to developing this state-of-the-art solution. IBM takes pride in advancing our capabilities based on our customers' needs. We have continuously updated, enhanced and improved Interact over its 10+ year existence to become the most robust solution set available for real-time marketing. And IBM is committed to continuing to advance this offering through industry-leading marketing expertise, our customers' vast experience, and using the largest scientific research staff available to any company worldwide.

- **Uses a modified Bayesian model:** IBM has made several adjustments to the base Bayesian theorem to account for typical problems found in learning algorithms.
 - **Prevent overwhelming bias:** To prevent overwhelming bias the Interact learning models performs a confidence test that tests the statistical relevance in every model variable before that variable is used to predict the outcome.
 - **Time-weighted response:** The Interact models are designed to react to changes in customer preferences over time. Interact weighs more recent responses heavier than it weighs responses that have happened in the distant past.
 - **Entropy analysis:** Interact will look at each model variable and assess its likelihood to be predictive. Variables that prove to be non-predictive are removed from the model. This both speeds the processing of the use models at real time and removes noise that may lessen the models accuracy.
 - **Multiple models and flexibility of learning variables:** Interact can support any number of learning models running simultaneously. The marketer can choose the starting set of learning attributes to use for a specific set of offers and can extend the list of attributes at any time.
 - **Learning measurement and reporting:** The application has built in control logic that will test offers with and without learning engaged so that Interact can accurately estimate the lift gain from using the learning components. In addition, the system will report back on which attributes are predictive for each offer in the system.
 - **Reacting to the value as well as the acceptance likelihood:** While the base algorithm predicts the likelihood of response, it is designed to take in a multiplicative factor that can represent the value of the offer to the customer or organization. This way the selection of which offer is best can be based on both value and acceptance likelihood.
- **Segment-based control:** Although Interact's learning capabilities are "state-of-the-art", marketers often want their own logic and strategies to supersede or influence the outcome of which offers are selected for each customer and each in each marketing situation. The Interact engine enable marketers to control when learning is engaged and the scope of influence the learning models have.
- **Designed to integrate with any third-party modeling engine:** The out-of-the-box functionality for learning is Bayesian with the additional capabilities described above. The system is not limited to using only our internal models; models, including real-time models, from other systems can be included into the arbitration logic of Interact.
- **Real time:** Statistics regarding offer acceptance are gathered in real time across any number of customer characteristics and behavioral situations.

Conclusion

With IBM Interact, today's marketers can combat the desensitization of today's customers to outbound marketing by communicating with them on their own terms—when they are ready and willing to engage. Through Interact's powerful, behavioral targeting analytics and marketer-defined business logic, marketers can deliver the optimal marketing message in each inbound channel—increasing targeted marketing opportunities, improving response rates in inbound channels, enhancing the customer experience, leveraging investments in inbound channels and maintaining a consistent dialogue with customers across all channels.

About IBM Enterprise Marketing Management

The IBM Enterprise Marketing Management (EMM) Suite is an end-to-end, integrated set of capabilities designed exclusively for the needs of marketing and related organizations. Integrating and streamlining all aspects of marketing, IBM's EMM Suite empowers organizations and individuals to turn their passion for marketing into valuable customer relationships and more profitable, efficient, timely, and measurable business outcomes.

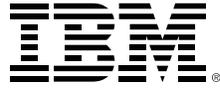
Delivered on premises or in the Cloud, the IBM EMM Suite of software solutions gives marketers the tools and insight they need to create individual customer value at every touch. The IBM EMM Suite helps marketers to understand customer wants and needs and leverage that understanding to engage buyers in highly relevant, interactive dialogs across digital, social, and traditional marketing channels.

Designed to address the specific needs of particular marketing and merchandising users, the IBM EMM Suite is comprised of five individual solutions. Digital Marketing Optimization enables digital marketers to orchestrate relevant digital interactions to attract and retain new visitors and grow revenue throughout the customer's lifecycle. With Customer Experience Optimization eCommerce professionals can turn visitors into repeat customers and loyal advocates by improving the digital experience of every customer. With Cross-Channel Marketing Optimization customer relationship marketers can engage customers in a one-to-one dialogue across channels to grow revenue throughout the customer's lifecycle. Price, Promotion and Product Mix Optimization allows merchandisers and sales planners to make price, promotion and product mix decisions that maximize profit and inventory utilization. And with Marketing Performance Optimization, marketing leaders, planners and decision-makers can model and assess mix, and manage marketing operations to maximize ROI.

Over 2,500 organizations around the world use IBM EMM solutions to help manage the pressures of increasing marketing complexity while delivering improved revenue and measurable results. IBM's time-tested and comprehensive offerings are giving companies such as Dannon, E*TRADE, ING, Orvis, PETCO, Telefonica | Vivo, United Airlines and wehkamp.nl the power and flexibility required to provide their customers and prospects with what they expect today—a more consistent and relevant experience across all channels.

For more information

To learn more about IBM Interact and IBM Enterprise Marketing Management please contact your IBM marketing representative or IBM Business Partner, or visit the following website: ibm.com/software/marketing-solutions.



© Copyright IBM Corporation 2013

IBM Corporation
Software Group (or appropriate division)
Route 100
Somers, NY 10589

Produced in the United States of America
March 2013

IBM, the IBM logo and ibm.com are trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the Web at “Copyright and trademark information” at www.ibm.com/legal/copytrade.shtml.

This document is current as of the initial date of publication and may be changed by IBM at any time. Not all offerings are available in every country in which IBM operates.

THE INFORMATION IN THIS DOCUMENT IS PROVIDED “AS IS” WITHOUT ANY WARRANTY, EXPRESS OR IMPLIED, INCLUDING WITHOUT ANY WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND ANY WARRANTY OR CONDITION OF NON-INFRINGEMENT. IBM products are warranted according to the terms and conditions of the agreements under which they are provided.



Please Recycle
