CHAPTER 13

Real-Time Inbound Marketing: A Use Case for Digital Banking

Alan Megargel, Venky Shankararaman, Srinivas K. Reddy

Contents

13.1 Introduction

In this era of digital banking, customers are becoming increasingly more sophisticated, expecting personalized banking services to be delivered to them anytime, anywhere, and across any channel. Customers are now unresponsive to banks' direct marketing tactics that target whole demographic segments rather than to them individually. Rather customers expect banks to know their situational needs, predict a next best offer, and deliver that offer to them at the right time and place.

Let's say you are shopping for cologne for yourself. Your credit card is swiped, and 2 seconds later you receive an SMS from your bank saying "we know it's your wife's birthday in 2 days, and we know she likes Gucci products, and we offer 15% discount on Gucci if you use our credit card, and there is a Gucci shop right across the corridor from you" (Ranadivé and Bilski, 2013). Your first reaction might be that you are thankful for the reminder because you indeed forgot about your wife's upcoming birthday, and your second reaction might be

Handbook of Blockchain, Digital Finance, and Inclusion, Volume 1
DOI: 10.1016/B978-0-12-810441-5.00013-0

to go across the corridor to Gucci and buy something for your wife. Furthermore, you would be delighted that your bank really knows you. The offer was delivered to you at precisely the right time, because you were just leaving one shop having purchased your cologne. If you had received the SMS 20 minutes later, already in your car, it would have been too late.

So how would a bank accomplish such a feat? In fact, only a few banks do it well, and many other banks are trying to follow their lead. This capability to deliver real-time personalized offers, at the customer point of interaction with the bank or merchant, is becoming increasingly strategic. Banks that fall behind will certainly lose market share. For the solution, you will need to look deep into the bank's architecture at the enterprise platform layer. There you will find enterprise platforms such as: messaging middleware, process orchestration engines, in-memory data grids, and complex event processing (CEP) engines. These are the architectural building blocks that make real-time inbound marketing happen.

This chapter proposes the usage of real-time inbound marketing as an important use case for a digital bank. The chapter will first discuss the evolution of marketing in retail banking along with the IT systems that support it, followed by an analysis of marketing strategies in today's digital banking era. The chapter will then provide a conceptual view of the various technology enablers of the enterprise platform layer of a bank's architecture. Finally, it will present a novel set of CEP patterns for implementing real-time inbound marketing, along with business scenario use cases for each pattern.

13.2 Evolution of IT Systems and Marketing Methods in Retail Banking

Marketing methods in retail banking are highly dependent on IT systems capability. The evolution of marketing methods can therefore be best understood in relation to the evolution of IT systems in retail banking. The evolution of IT systems can be explained in three separate eras: the "data processing era", the "client–server era", and the "predictive era". As IT systems evolved over these three eras, marketing methods evolved from batch mode outbound marketing to mass consumers without segmentation, followed by outbound marketing with customer segmentation, and finally to real-time personalized inbound marketing to individual consumers.

Data Processing Era. The "data processing era" (or "mainframe era") took place in the 1960s through 1980s. In this era, banking transactions were typically made over the counter, during normal office hours – 8 by 5. A typical bank with 10 million customers, for example, would interact with customers in person at the branch, and the number of customer interactions per day could be measured in the thousands.

A transaction was the only record of an interaction with a customer. A business day's transactions were processed overnight in batch mode (Martin, 2012), which meant that any potential marketing opportunities associated with customer transactions could only be assessed on the next day. In this era, customers had little or no access to information technology, and therefore products and services offered to customers were limited and unsophisticated relative to today. Furthermore, banking systems in this era were focused around accounting (Martin, 2012), and did not accommodate customer segmentation for the purposes of marketing. Banking was mostly relationship based, and with a relatively low level of sophistication, the speed of doing business was slow.

Client–Server Era. The "client–server era" took place in the 1980s through 2000s. The client–server architecture was meant to enable a bank's IT infrastructure to scale more cost effectively as compared to the mainframe, and it enabled more users to access their business systems. "Thick" client applications led to "thin" client web applications, and products and services were then offered over the Internet. With the advent of the Internet and other self-service channels, customers became more sophisticated and demanded better products and faster service, 24 by 7. As a result, the speed of doing business increased, in order to stay competitive. The number of customer interactions per day could be measured in the millions.

Self-service channels like the Internet resulted in more customer touch points, but these events did not necessarily get captured as a transaction. For example, if a customer logged in and browsed a bank's web site and showed interest in a product or service (e.g., used a simple mortgage calculator), no transaction took place, and therefore there was no record of the customer's interest.

In this era, banking operations were still transaction based. A transaction was still the only record of an interaction with a customer. But with the speed of business increasing, more transactions were being stored. Relational databases became popular for operational systems that processed a day's transactions. Enterprise data warehouses (EDWs) became popular to store historical transactions, but as a carryover from batch-style thinking, EDWs would only be updated during end-of-day batch processing.

To gain competitive advantage, banks began to mine their EDWs to generate marketing leads based on historical customer transaction behavior. Business Intelligence (BI) software became popular to help segment customers, and to identify opportunities (e.g., cross-sell, up-sell) as well as threats (e.g., late payments, fraud). However, an EDW-based BI could only access transaction data that was one day old at best, and therefore real-time business decisions were not possible due to "information latency" (TIBCO, 2006a).

In the client–server era, it was not possible for a bank to predict in real-time what products or services a customer might be interested in. Nor was it possible for a bank to predict in real-time if they were going to lose a customer. This era could also be described as the "wish I knew era," because once a critical business event had passed then it was too late to do anything about it.

Predictive Era. We are currently in the "predictive era" or the "we know what's about to happen era" (Ranadivé, 2005). The speed of doing business demands that events from multiple sources be captured and correlated in real-time, so that the bank can know ahead of time what is about to happen, and do something about it before it is too late. In the previous era, banking operations were transaction based, and decisions were made based on historical data-at-rest sitting in a database or EDW. In the current era, decisions are also made based on data-in-motion which can be either transactions or business events that are evaluated in real-time, in memory. To a predictive bank, data-in-motion is an order of magnitude more valuable than data-at-rest (Ranadivé and Maney, 2011).

Business Events can originate from a multitude of different sources, including: customer interactions (self-service and assisted), process milestones, system notifications, etc. Events can be temporal in nature also: for example, if something occurs earlier or later than expected, then that is useful information. Or if something doesn't occur at all, then that is an event (e.g., if a person starts an online loan application, but does not submit it). When assessed in isolation, a single event might not be enough information to predict something is about to happen. However, when events from different sources are correlated, patterns can be formed that lead to a prediction. For example, if a person is late renewing their credit card membership, and this person logs a complaint with the bank's call center, and within a week this person transfers a large sum of money to another bank, then a prediction can be made that the bank will lose this customer unless they do something about it.

When customer interaction events, such as: a credit card swipe, a web page visit, an ATM transaction, a bill payment, etc., are correlated with life-stage events, such as: getting married, moving house, having children, retiring, etc. (Bailey et al., 2009), there arise opportunities to effectively cross-sell or up-sell specific products targeting specific customers in real-time at the point of customer interaction with the bank (TIBCO, 2006b). Decisions made in real-time on these opportunities enable banks to intelligently sell the right products to the right customers via the right channels (Kamakura, 2008). The real-time execution of a perfectly customized and personalized "next best offer" (NBO) is the holy grail of inbound marketing (Davenport et al., 2011). Banks that do this well will have a competitive advantage.

Table 13.1 summarizes the key attributes for each era.

Table 13.1: Comparison of the different era's.

	Period	Customer Interaction	Marketing Strategy	IT Architecture	Transaction Processing
Data Processing Era	1960s–1980s	Face-to-Face over the counter. Around 1000 interactions/day through face-to-face over the counter.	Relationship based	Internally focused infrastructure with mostly the bank officers using systems	End of day batch processing of transactions
Client–Server Era	1980s–2000s	Many touch points for customer interaction. However not all interactions are captured as transactions and correlated.	Marketing leads based on historical customer data and segmentation	Internally and externally focused infrastructure with both bank officers and customers accessing services over the internet 24 by 7	Transactions in real-time but analytics and decision making are not real-time
Predictive Era	Now	Many touch points for customer interaction. All interactions are captured as transactions and correlated in real-time	Marketing leads based on historical and data in motion, segmentation, current context	24 by 7 open but secure infrastructure across the business ecosystem of bank, partners and customers	Transactions, analytics and decision making are in real-time

13.3 Marketing Strategies in Today's Digital Banking Era

Digital Banking. "A digital bank uses a broad range of technology-centric capabilities that enable new methods of interaction and service delivery to augment the customer experience and potentially transform the business. These capabilities are supported by a robust, dynamic and accessible digital infrastructure and open banking system that transform the analog environment" (Gartner Research, 2013b).

In the digital banking market, retail banks can no longer compete solely based on products which have become commoditized, rather they must compete based on customer experience driven service differentiation (Drotskie, 2009). Digital natives, who have been immersed in technology their entire lives (Prensky, 2001), are unresponsive and even annoyed with mass marketing promotions that are irrelevant to them personally, rather they expect technologically sophisticated and personalized treatment (Cotton and Walker, 2011; Drotskie, 2009).

Cross-Selling Strategy. Brand loyalty based on relationships is not sufficient. As banking customers are exposed to increasingly more innovative self-service channels, these now sophisticated customers tend to shop around and do business with every bank that offers the best service, which usually means the fastest delivery (Cotton and Walker, 2011). Cross-selling products to existing banking customers is more cost effective, roughly one-fifth the cost of selling to new customers (Kona and Surti, 2010), and also improves customer loyalty as the number of products per customer increases. Cross-selling therefore, as a means to both increase revenue and retain customers, has become a key strategic initiative for many retail banks (Dass, 2006; Kamakura, 2008; Mann and Kumar, 2014).

Traditionally, cross-selling required person-to-person interaction at the branch and relied on the skill and intuition of the bank officer. However, in the digital banking market, self-service online channels have largely supplanted human interactions at the branch, and therefore IT systems and analytics tools are now central to cross-selling (Kamakura, 2008).

Cross-selling has implications on a bank's channel integration strategy, as increasingly more customers do their banking across multiple channels, and expect consistent treatment across all channels (Cotton and Walker, 2011; Stone, 2009). The effectiveness of product offers delivered across multiple channels needs to be understood and optimized, and marketing content needs to be designed and optimized specifically for each channel. Customer interaction management is needed as well which: a) enables customers to opt-in to receive promotions and select their preferred mode of delivery, and b) enables the bank to track whether or not a customer has expressed interest in the offer in order to avoid irritating the customer with repeated offers for a product they are not interested in.

Cross-sell offers can be delivered via assisted channels such as branch and call center, provided that the branch teller or call center operator is prompted with an auto-generated "Next Best Offer" (NBO) tailored for the customer (Hesse, 2009). This is typically the first form of inbound marketing capability implemented by banks, although the reach to customers is limited as increasingly more customers prefer to do their banking through online channels. Still, it is an important milestone for a bank to achieve, since the capability to deliver tailored offers across multiple channels implies that the bank has a 360-degree view of customers supported by predictive analytics (Cognizant, 2011).

In the current predictive era, a bank's inbound marketing strategy must include online channels and payments as the predominant sources of customer interaction. In Asia, 40% of mass affluent customers and 50% of customers under the age of 40 prefer to use Internet and mobile banking channels (McKinsey & Company, 2014). In the US, 80% of all customer interactions with the bank are for payment-related activities such as making payments, checking

payment status, or paying bills (Denecker et al., 2014). Therefore, banks must invest in technology that enables cross-sell offers to be delivered via online channels in real-time, i.e., at the moment of customer interaction with the bank.

Simple Message Service (SMS) text messaging is a cost-effective and reliable channel for delivering targeted cross-sell offers to customers at the point-of-sales (POS) in real-time just after a credit card swipe (Riley et al., 2011). If the merchant has a tie-in with the bank, the cross-sell offer might be related to that merchant and offer a discount coupon. The merchant may not have a tie-in with the bank however, and the customer interaction might then be used to cross-sell another nearby merchant's offer which is relevant to the customer. Either way, it is a NBO that is delivered in real-time.

A Framework for NBO. Davenport et al. (2011) propose a four-step framework for enabling NBO, namely defining objectives, by gathering data, analyzing and executing, and learning and evolving. In the defining objective step, the bank must define what it intends to achieve as the end goal, for example, increased revenues, increased customer loyalty, increased customer purchase or new customers. In the gathering data step, the bank must collect and integrate detailed data about their customers, product and service offerings, and the context in which customer purchases the products. In the analyzing and executing step, the bank must identify the right technology for analyzing and predicting the offer to be offered, and also the channel (e.g. call center, online in real-time) through which the NBO is to be delivered. In the learning and evolving step, the bank must put in place mechanisms to monitor and learn from their NBO performance and evolve their strategy and implementation.

Organizational Challenges. Cross-selling effectively requires a change in marketing method; from traditional outbound campaign-driven mass marketing based on customer segmentation, to inbound marketing which delivers personalized offers targeted to individual customers based on their situational needs. This presents several organizational challenges:

Firstly, traditional outbound marketing is very much product-centric, aligning to traditional product-centric organizational silos (Kamakura, 2008). A marketer designs an offer for a banking product, selects a segment of customer's to target, and then delivers the same offer through the same channel to all customers in the selected segment (Hesse, 2009). Banks will need to unlearn what they have learned in order to break down their product-centric mentality. Inbound marketing requires a customer-centric mindset (Kamakura, 2008). A customer interacts with the bank using their preferred channel, the customer's profile and current situation determine which one of many possible product offers is best suited for the customer at that point in time, and the offer is delivered in real-time via the customer's preferred delivery channel.

Secondly, there may be management resistance to implementing a customer-centric inbound marketing strategy, due to existing sales incentives which are traditionally tied to product performance (Kamakura, 2008). New creative incentive schemes are needed, possibly tied to customer value or cross-sell related metrics.

Thirdly, banks may lack the right mixture of expertise. Inbound marketing requires a combination of: data analytics expertise, banking domain knowledge, and decision support systems expertise. Marketing activities would need to be supported by a cross-functional team of business users having product knowledge and customer experience knowledge, as well as data analysts skilled in analyzing and interpreting large amounts of data from various sources (Kona and Surti, 2010). Actionable cross-sell decisions in the form of business rules would then be deployed by decision support systems specialists.

Citibank Example. The capability to deliver real-time personalized offers, at the customer point of interaction with the bank or merchant, is becoming increasingly strategic. Only a few banks do it well, and many other banks are trying to follow their lead. One example of a successful real-time inbound marketing implementation is at Citibank, where they have developed a real-time marketing engine they call Centralized Offer Pallet System (COPS).

"COPS is able to instantaneously consider a number of customer events, including banking events, such as credit card purchases or ATM transactions, and life events, such as customer birthdays or overseas travel, when determining the bank's response. The system deploys complex event processing technology to evaluate static and dynamic events against a customer profile and 'propensity model', to determine in real-time the next best offer the bank can extend to the customer" (The Banker, 2009).

13.4 Technology Enablers for Real-Time Inbound Marketing NBO

It the present predictive era, the speed of doing business demands real-time decision making. Predictive capability requires massively scalable highly resilient event-driven architecture, complex event processing, and high performance in-memory analytics. Massive scale, in some cases, means millions of events per second originating from a multitude of different event sources need to be collected and correlated in real-time, so that the bank can know ahead of time what is about to happen, and do something about it before it is too late to take advantage of an opportunity or avert a threat.

With massive amounts of events being captured, data volumes are increasing (see Fig. 13.1). However, the time to react to opportunities or threats is decreasing. After it is too late to react, the data becomes less useful (Fülöp et al., 2012, September). At the same time, the cost of solid state storage and the cost of memory are decreasing rapidly. While events (data-in-motion) are becoming business critical, databases (data-at-rest) are becoming less relevant

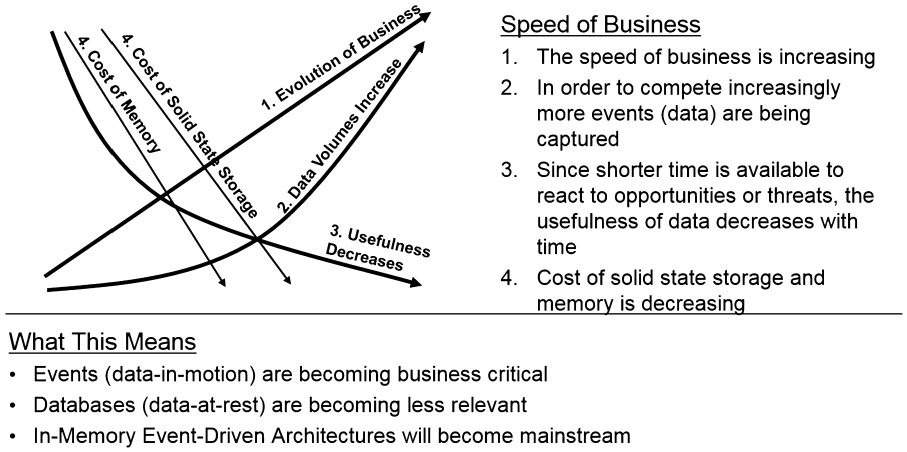

Figure 13.1: Shift towards In-Memory Event-Driven Architectures.

(Ranadivé and Maney, 2011). Databases will always be required to store historical transaction data. But as the cost of memory is decreasing, massive scale high performance in-memory analytics and real-time decision systems will become mainstream.

An Event-Driven Architecture (EDA) is one in which business events, indicating a change in state, are detected and automatically trigger business processes to activate (Shankararaman and Megargel, 2013). For example, a creditworthy customer changing their mailing address online might trigger a business process to cross-sell a furniture promotion, whereby the furniture store has a merchant tie-in with the bank. The essential enterprise platform components of an EDA, given in the typical order of adoption by banks, are as follows:

Message-Oriented Middleware. A means by which applications share information across the network, message-oriented middleware (MOM) is typically the first integration technology adopted by banks. Middleware vendors offer MOM software that is compliant to the Java Message Service (JMS) standard as ratified by the World Wide Web Consortium (W3C). JMS is preferred by banks as it is robust, secure, high performing, and guarantees delivery of messages which often contain customer sensitive financial information. In an EDA, JMS messages are an "event source" which trigger real-time inbound marketing offers.

Service-Oriented Standards. Services are reusable software components that expose functionality via a defined interface, for example, a balance inquiry service. An enterprise service bus (ESB) exposes the functionality of an enterprise as reusable services. A service oriented architecture (SOA) enables new applications to be "assembled" rapidly by invoking existing reusable services via an ESB. The rapid construction and assembly of reusable services requires a high degree of consistency in interface design, which is why most banks comply with

W3C standards for SOA including, Simple Object Access Protocol (SOAP), and Web Services Description Language (WSDL). As a bank's level of SOA maturity increases, so does its business agility.

Established banks that are still using monolithic legacy core banking systems, which are inflexible to change, have a higher barrier to entry into the digital banking market currently lead by FinTech and other IT companies. "Legacy core banking platforms across the industry, traditionally dependent on mainframes, are now giving way to platforms based on service-oriented architectures" (Cognizant, 2011).

Business Process Orchestration. Business Process Management (BPM) platforms are used by banks to model, simulate, execute and monitor complex business processes such as customer on-boarding, credit evaluation, and loan origination. BPM platforms orchestrate business processes that include human steps as well as machine steps. For each machine step, the BPM platform will invoke a reusable service via an ESB, based on the WSDL that defines the interface of that service. In an EDA, a JMS message as an event source might activate a BPM process which orchestrates an inbound marketing offer to the customer.

Complex Event Processing. "Complex-event processing (CEP) is a kind of computing in which incoming data about events is distilled into more useful, higher level 'complex' event data that provides insight into what is happening. CEP is event-driven because the computation is triggered by the receipt of event data. CEP is used for highly demanding, continuous-intelligence applications that enhance situation awareness and support real time decisions" (Gartner Research, 2013a). A mid-size bank might have thousands of simple events occurring every second, forming an "event cloud" (Luckham, 2008), consisting of deposits or withdrawals, fund transfers or bill payments, credit card swipes or online purchases.

At the core of a CEP is a rule engine capable of composing context-sensitive complex events (or rules) out of the thousands of simple events that occur every second (Adi et al., 2006). Context-sensitive rules encompass: historical customer banking behavior and credit worthiness, elapsed time between related events, proximity or location of events, as well as algorithmic measurement or aggregation of numeric event data (Adi et al., 2006; Luckham, 2008; Lundberg, 2006). Rules are externalized or decoupled from the business process execution such that the rules are independently managed and reusable across multiple business processes (Adi et al., 2006).

In a bank, there are typically two business units which are involved in setting up a real-time inbound marketing NBO. A decision support unit will use a CEP designer tool to: a) model the context-sensitive threshold, variance, temporal and spatial constraints, algorithms and data aggregations which constitute a complex event (or rule), and b) deploy the rule onto a CEP

runtime execution environment (TIBCO, 2006a). A marketing team will then: a) design a promotion for a banking product or merchant product associated with the bank, and b) use the CEP designer tool to select a set of rules which will trigger the promotion to be delivered to targeted customers whom have opted-in to receive such promotions. The NBO will then be deployed onto an in-memory data grid, queued up to be delivered in real-time to each individual targeted customer upon their next interaction with the bank (TIBCO, 2006b), ergo the NBO delivered in real-time.

See the Appendix for CEP patterns and use cases for real-time inbound marketing.

The term "real-time" in the context of inbound marketing means roughly "less than 2 seconds" as a key performance benchmark, which means that the targeted customer will receive a relevant NBO within 2 seconds of their next interaction with the bank, typically via SMS. If NBOs were to be retrieved from a customer analytics data mart situated downstream from a large data warehouse, the overhead of "seek" access-times while reading data from physical disk drives would inhibit the 2 second performance benchmark. Disk drive access times are measured in milliseconds, whereas memory access times are measured in nanoseconds. Therefore, real-time NBOs are more effective when stored in memory rather than on disk.

From 1956–2015, the cost per megabyte of disk drives has fallen from $9,200.00 to $0.0000317 (McCallum, 2015a). From 1957 to 2015, the cost per megabyte of random access memory (RAM) has fallen from $411,014,792 to $0.0056 (McCallum, 2015b). The price of disk drives has leveled off since 2006, whereas the price of RAM has halved since 2011 (McCallum, 2015a, 2015b). The cost of RAM is decreasing faster than the cost of disk drives.

In-Memory Data Grid (IMDG). "In-memory data grids are distributed, in-memory data stores aimed at high performance/high-scale, data-intensive applications. IMDGs are moving into mainstream adoption, driven by their versatility, including support for cloud architectures, and bundling in various software products and services" (Gartner Research, 2012). IMDGs provide a virtual shared memory by replicating data elastically across any number of active–active fault tolerant nodes (servers) in a peer-to-peer style of architecture. IMDGs store and retrieve in-memory data using key-value pairs, and can "push" or "listen" to events in support of event driven architectures and CEP. IMDGs are therefore well placed to store millions of NBOs in memory, queued up to deliver inbound marketing offers in real-time. Memory access-times are around a million times faster than disk drive access-times, and the low cost of memory is catching up with the lower cost of disk drives. In the context of real-time inbound marketing, databases are becoming less relevant. Hence, there is a shift towards in-memory event driven architectures. "Memory is the new disk, disk is the new tape", coined by Jim Gray (Robbins, 2008).

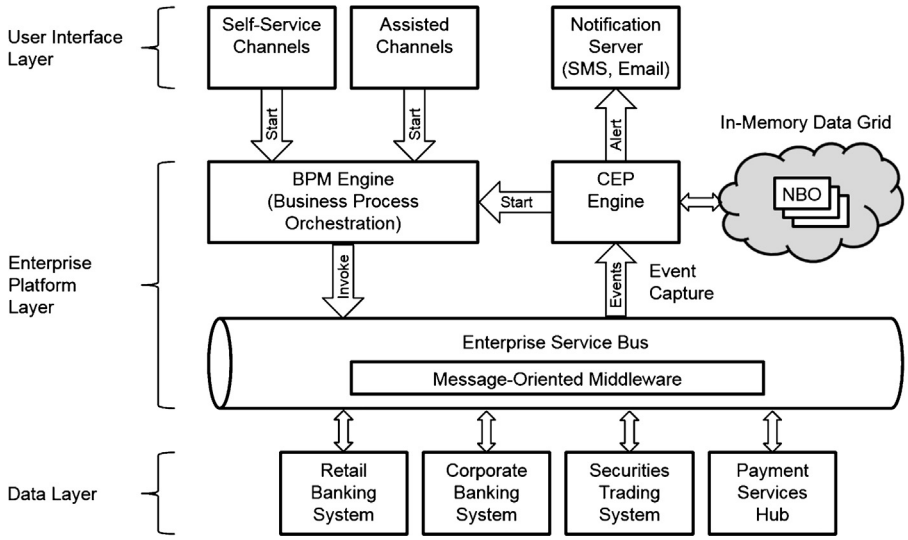

Figure 13.2: Layered Architecture to support Real-time In-Bound Marketing.

Layered Architecture for Real-time Inbound Marketing. The various technology components collaborate to support real-time inbound marketing. Fig. 13.2 shows the interactions between the technology components. In order to understand the interactions, let us use the following CEP Pattern 4: Life-stage event – merchant product cross-sell (use case 4.3 from the appendix):

"Customer John Tan is moving house. He logs into bank's website and changes mailing address."

The customer initiates the "address change" service through the self-service channel. This triggers the corresponding business process in the BPM engine, and through the enterprise service bus (ESB) the appropriate service, for example, "change_address" is invoked in the Retail Banking System. Simultaneously, this business event is captured by the complex event processing (CEP) engine. Within the CEP engine, the following marketing business rule is triggered:

"IF business event = "change_address" THEN send "furniture_cross_sell SMS to customer"

From the In-Memory Data Grid data, an affiliated furniture store nearest to the customer's new mailing address is selected, and a SMS is sent with the following data"

"Dear Mr John Tan we have updated your address in the bank system. We are glad to offer you a credit card discount of 15% for purchases above $200 at the Myfurniture store located at 20 Queen Street".

13.5 Summary

This chapter first examined the evolution of marketing methods in retail banking, as a function of IT systems capability across three separate eras; the "data processing era", the "client–server era", and the "predictive era". The chapter then discussed banks' marketing strategy in the context of today's digital banking market, with emphasis on cross-selling and its related benefits and organizational challenges. Finally, the chapter covered the banking industry shift towards in-memory event-driven architectures which enable the real-time delivery of personalized "next best offers" targeting individual customers based on their profile and situational needs.

In the digital banking market, retail banks can no longer compete solely based on products which have become commoditized, rather they must compete based on service differentiation. Digital natives who expect technologically sophisticated and personalized service are unresponsive to traditional campaign-based outbound marketing methods. Inbound marketing processes implemented at assisted channels such as branch and call center have a limited reach as increasingly more customers prefer to do their banking through online channels. Therefore, banks must invest in technology that enables cross-sell offers to be delivered via online channels in real-time, i.e. at the moment of customer interaction with the bank. Banks that do this well will have a competitive advantage. Banks that do not adopt this new paradigm will not survive.

Appendix – CEP Patterns and Use Cases for Real-Time Inbound Marketing

Assumptions for the below use cases:

- Customers have opted-in to receive real-time offers.
- Offers occur roughly within 2 seconds of the event, i.e. a customer interaction.
- Offers are not repeated, or the frequency of repeated offers is otherwise optimized.
- "[NBO]" indicates offer is queued up in memory, to be triggered at the next customer interaction.

See Tables 13.2–13.6.

Table 13.2: CEP Pattern 1.

CEP Pattern 1: Life-stage event – banking product cross-sell		
Use Case	**Event**	**Offer**
1.1	Customer enters university. Pays application fee using the bank's credit card.	An SMS is sent to the customer with education loan rates, a link to the bank's web page, 3 possible appointment times at the branch nearest the customer's address, and the phone number of a relationship manager (RM).
1.2	Customer's firstjob. First payroll direct deposit for a customer having no credit card.	[NBO] An SMS is sent to the customer with credit card rates, and a link to the bank's web site. A credit card application form is sent to the customer via regular mail.
1.3	Customer is buying a house. Logs into bank's website and uses a mortgage calculator.	A robo-chat box pops up providing credit terms computed based on the customer's credit worthiness. Dialog leads to connection with a human to facilitate loan application.
1.4	Customer has first child. Pays for a crib or high chair using the banks credit card.	An SMS is sent to the customer offering a "Baby Bonus" savings plan. A loan application form is sent to the customer via regular mail.
1.5	Customer reaches middle age.	[NBO] An SMS is sent to the customer offering wealth management advisory services, and the phone number of an RM. A wealth management brochure is sent to the customer via regular mail.
1.6	Customer reaches retirement age.	[NBO] An SMS is sent to the customer offering an annuity based on the customers accumulated wealth, and the phone number of an RM. An annuity plan is sent to the customer via regular mail.

Table 13.3: CEP Pattern 2.

CEP Pattern 2: Financial event – banking product cross-sell		
Use Case	**Event**	**Offer**
2.1	Dormant account. Customer has a large balance in a demand deposit account which has been dormant for an extended period.	[NBO] An SMS and/or email is sent to the customer offering a higher yielding structured deposit product. The dormant account is mentioned in the communications, and the phone number of an RM is provided.
2.2	Mortgage full repayment inquiry. Customer uses the bank's website to compute the full repayment of existing mortgage.	A robo-chat box pops up to inquire if the customer is planning to refinance with another bank, or is planning to sell and reinvest in another property. Dialog leads to connection with a human to offer competitive rates.
2.3	Direct debit authorization with a competitor. Insurance company is authorized to debit your customer's account directly.	[NBO] An SMS and/or email is sent to the customer to inform them of your bank assurance products that they may not be aware of. A link to the bank's relevant web page is included in the communications.
2.4	Salary bonus. Customer's direct payroll deposit is much larger than normal.	[NBO] An SMS and/or email is sent to the customer mentioning their bonus and offering an investment product featuring a lump sum premium based on the customer's value/portfolio with the bank.
2.5	Large credit card transaction. Customer who is a "transactor" (not a "revolver") makes a large credit card transaction.	An SMS is sent to the customer offering to convert their credit card debt to a lower interest personal loan, providing rates computed according to the customer's credit worthiness, and the phone number of an RM.
2.6	Time Deposit renewal. A customer's TD maturity is approaching.	[NBO] An SMS and/or email is sent to the customer offering a renewal of their TD, with more attractive rates if they increase the deposit amount.
2.7	Flight Insurance. Customer purchases a budget airline ticket using the bank's credit card, while at the airport.	An SMS is sent to the customer offering cheap flight insurance, and a link to the bank's relevant web page where they can purchase the insurance online.
2.8	ATM insufficient balance. Customer attempts an ATM withdrawal but has insufficient funds in their account.	An SMS is sent to the customer offering a short term personal loan, with rates based on the customer's credit worthiness, and a link to the bank's relevant web page to apply for the loan online.

Table 13.4: CEP Pattern 3.

CEP Pattern 3: Point-of-sales (POS) event – merchant product cross-sell		
Use Case	**Event**	**Offer**
3.1	Customer makes a POS purchase at an affiliated merchant, using the bank's credit card.	An SMS is sent to the customer offering a credit card discount for the next purchase at the same merchant, redeemable one time only via a QR code. The customer receives merchant loyalty points.
3.2	Customer makes a POS purchase within walking distance of an affiliated merchant where the customer has an affinity.	An SMS is sent to the customer offering a credit card discount at the nearby affiliated merchant. The bank establishes customer affinity for the merchant at the originating POS.

Table 13.5: CEP Pattern 4.

CEP Pattern 4: Life-stage event – merchant product cross-sell		
Use Case	**Event**	**Offer**
4.1	Birthday or anniversary. Family member has an approaching birthday or anniversary.	[NBO] An SMS is sent to the customer offering a credit card discount at an affiliated merchant where the customer's family member has an affinity.
4.2	Customer changes jobs. The customer's direct payroll deposit has changed to originate from a new employer's account.	[NBO] An SMS is sent to the customer offering a credit card discount at an affiliated clothing store where the customer has an affinity.
4.3	Customer is moving house. Logs into bank's website and changes mailing address.	An SMS and/or email is sent to the customer offering a credit card discount at an affiliated furniture store in the proximity of the customer's new mailing address.
4.4	Customer is taking a trip. Books a hotel using the bank's credit card.	An SMS and/or email is sent to the customer offering a credit card discount with an affiliated airline, where the loyalty programs of the bank and airline are linked.

Table 13.6: CEP Pattern 5.

CEP Pattern 5: Location-based event – merchant product cross-sell		
Use Case	**Event**	**Offer**
5.1	Customer retrieves ATM card from the machine, and the time of day is around 11am or 5pm.	An SMS is sent to the customer offering a credit card discount at an affiliated restaurant where the customer has an affinity, and is within walking distance of the ATM.
5.2	GPS tracking is enabled on the customer's mobile phone, and the customer is on foot in a downtown area.	An SMS is sent to the customer offering a credit card discount at an affiliated merchant where the customer has an affinity, and the merchant is within walking distance of the customer's current location.

References

Adi, A., Botzer, D., Nechushtai, G., Sharon, G., 2006. Complex Event Processing for Financial Services. In null. IEEE, pp. 7–12.

Bailey, C., Baines, P.R., Wilson, H., Clark, M., 2009. Segmentation and customer insight in contemporary services marketing practice: why grouping customers is no longer enough. Journal of Marketing Management 25 (3–4), 227–252.

Cognizant, 2011. Cognizant reports: how analytics can transform the U.S. retail banking sector. Retrieved December 15, 2015, from http://www.cognizant.com/InsightsWhitepapers/How-Analytics-Can-Transform-the-U.S.-Retail-Banking-Sector.pdf.

Cotton, B., Walker, C., 2011. Smarter customers, smarter commerce: innovations for the banking industry.

Davenport, T.H., Mule, L.D., Lucker, J., 2011. Know what your customers want before they do. Harvard Business Review 89 (12), 84–92.

Dass, R., 2006. Data Mining in Banking and Finance: A Note for Bankers. Indian Institute of Management Ahmedabad.

Denecker, O., Gulati, S., Niederkorn, M., 2014. The Digital Battle that Banks Must Win. McKinsey & Company, Washington, DC. http://www.mckinsey.com/insights/financial_services/the_digital_battle_that_banks_must_win.

Drotskie, A., 2009. Customer Experience as the Strategic Differentiator in Retail Banking. Doctoral dissertation. University of Stellenbosch, Stellenbosch.

Fülöp, L.J., Beszédes, Á., Tóth, G., Demeter, H., Vidács, L., Farkas, L., 2012, September. Predictive complex event processing: a conceptual framework for combining complex event processing and predictive analytics. In: Proceedings of the Fifth Balkan Conference in Informatics. ACM, pp. 26–31.

Gartner Research, 2012. What IT leaders need to know about in-memory data grids. Publication G00231619, March 6, 2012. Retrieved December 18, 2015, from https://www.gartner.com/doc/1942417/it-leaders-need-know-inmemory.

Gartner Research, 2013a. Gartner IT Glossary: Complex-Event Processing. Retrieved December 18, 2015, from http://www.gartner.com/it-glossary/complex-event-processing.

Gartner Research, 2013b. What is digital banking? Publication G00255405, September 30, 2013.

Hesse, A., 2009. Case study: ING delivers personalized product offers across channels in real time.

Kamakura, W.A., 2008. Cross-selling: offering the right product to the right customer at the right time. Journal of Relationship Marketing 6 (3–4), 41–58.

Kona, Y.A., Surti, B.C., 2010. Analytics in cross selling – a retail banking perspective.

Luckham, D., 2008. Complex event processing in financial services.

Lundberg, A., 2006. Leverage complex event processing to improve operational performance. Business Intelligence Journal 11 (1), 55.

Mann, P.K., Kumar, R., 2014. Effectiveness of E-CRM in banking sector.

Martin, I., 2012. Too far ahead of its time: Barclays, Burroughs, and real-time banking. IEEE Annals of the History of Computing 34 (2), 5–19.

McCallum, J.C., 2015a. Disk drive prices (1955–2015). Retrieved December 18, 2015, from http://www.jcmit.com/diskprice.htm.

McCallum, J.C., 2015b. Memory prices (1957–2015). Retrieved December 18, 2015, from http://www.jcmit.com/memoryprice.htm.

McKinsey & Company, 2014. 2014 Digital banking in Asia – winning approaches in a new generation of financial services. Retrieved November 23, 2015, from http://www.google.com.sg/url?sa=t&rct=j&q=&esrc=s&source=web&cd=6&ved=0ahUKEwihjarL6aXJAhXi26YKHVuqDekQFghBMAU&url=http%3A%2F%2Fwww.mckinsey.com%2F~%2Fmedia%2Fmckinsey%2520offices%2Fsingapore%2F2014%2520digital%2520banking%2520in%2520asia%2520%2520winning%2520approaches%2520in%2520a%2520new%2520generation%2520of%2520financial%2520services.ashx&usg=AFQjCNFed4vE2ub6wyt2ajok5SidadAeBw.

Prensky, M., 2001. Digital natives, digital immigrants part 1. On the Horizon 9 (5), 1–6.
Ranadivé, V., 2005. The Power to Predict: How Real Time Businesses Anticipate Customer Needs, Create Opportunities, and Beat the Competition. McGraw-Hill Pub. Co.
Ranadivé, V., Maney, K., 2011. The Two-Second Advantage: How We Succeed by Anticipating the Future–Just Enough. Crown Business.
Ranadivé, V., Bilski, D., 2013. Real-time seer: Vivek Ranadivé. Outlook business, February 2013. Retrieved December 15, 2015, from http://blogcms.outlookindia.com/article_v3.aspx?artid=283868.
Riley, B., Schmidt, A., Tubin, G., 2011. SMS in financial services: accessing your customers on their terms. TowerGroup. Research is available on the Internet at www.towergroup.com.
Robbins, S., 2008. RAM is the new disk. . . . InfoQ News. Retrieved December 25, 2005, from http://www.infoq.com/news/2008/06/ram-is-disk.
Shankararaman, V., Megargel, A., 2013. Enterprise integration: architectural approaches. Service-Driven Approaches to Architecture and Enterprise Integration, 67.
Stone, M., 2009. Staying customer-focused and trusted: Web 2.0 and Customer 2.0 in financial services. Journal of Database Marketing & Customer Strategy Management 16 (2), 101–131.
The Banker, 2009. Innovation In Technology Awards 2009. The Chair's Choice & Innovation in Customer Service and Marketing Technology. Winner: Citibank Hong Kong. Project: Real-time marketing. The Banker, June 2009, p. 44.
TIBCO, 2006a. Complex event processing: framework for operational visibility and decisions.
TIBCO, 2006b. Predictive Customer Interaction Management: an architecture that enables organizations to leverage real-time events to accurately target products and services.