
Closing the CRM loop: The 21st century marketer's challenge: Transforming customer insight into customer value

Received (in revised form): 11th September, 2001

Anton Hirschowitz

is an experienced technical architect, focusing on CRM and customer insight solutions. He currently holds the position of principle CRM Consultant/Lead Technical Architect in the customer insight division of Detica, the largest specialist CRM consultancy in the UK.

Abstract Customer insight underpins customer relationship management (CRM). Without a detailed understanding of customer profiles and behaviour, any CRM undertaking will be running blind. Conversely, no matter how sophisticated a company's ability to generate customer insight, it will deliver little value without the processes in place that exploit this understanding to build stronger customer relationships. This paper illustrates how to deploy customer insight throughout the organisation to increase customer value and to learn from customer interactions to generate further customer insight. This 'closed-loop' approach is the crucial step necessary to bring the analytical and operational components of CRM into harmony.

This paper describes the types of techniques and technologies (such as campaign management) used to transform customer insight into customer value, and considers some of the pitfalls surrounding the use of these tools. The paper also outlines the best approach for integrating customer insight and customer interaction processes, including examples of best practice business and technical architecture solutions. Finally, it explains a 'recipe for success' for closed-loop CRM. This examines the key factors that organisations must consider most carefully when implementing the solution: customer targeting, response capture and performance measurement.

INTRODUCTION

Customer insight underpins customer relationship management (CRM). Without a detailed understanding of customer profiles and behaviour, any CRM undertaking will be running blind. Conversely, no matter how sophisticated a company's understanding of its customer, this insight will deliver little value without the processes in place that

exploit it to build stronger customer relationships.

As an organisation and its customers adapt to emerging e-channels, the challenge to develop and exploit customer insight will broaden and deepen as it attempts to encompass every class of marketing communication, from above-the-line campaigns down to real-time personalisation at the

Anton Hirschowitz
Detica, Surrey Research
Park, Guildford, Surrey GU2
7YP, UK.
Tel: +44 (0)1483 442000;
Fax: +44 (0)1483 442285;
e-mail: antonhirschowitz@
detica.com




Figure 1 'Closing the CRM loop'

interaction 'touchpoint' (call centre, website, third-generation mobile, etc).

This paper will explain the types of processes, techniques and tools needed to transform customer insight into increased customer lifetime value. By deploying customer insight into customer interactions, and understanding the responses to generate new insight, the vision of 'closed-loop CRM' is brought closer to reality.

WHAT IS CLOSING THE LOOP?

Closing the CRM loop involves three key processes:

- deploying customer insight into the operational CRM environment (through marketing, sales and service applications) to enhance the effectiveness of customer communications
- gathering data on marketing campaigns, channels, treatments and customer responses in order to provide visibility of the effectiveness of CRM activity

- analysing these data to generate customer insight, enabling analysts to understand how customers react to different forms and content of interaction.

In this way, it is possible to achieve full visibility of campaign and channel performance, and to exploit this knowledge to optimise the effectiveness of every customer contact.

APPLICATIONS

The main applications supported by customer insight in 'closed loop' CRM are:

- campaign management: a wide range of tools exists to take the leg work out of managing marketing campaigns, enabling companies to design and simultaneously manage large numbers of small campaigns that are accurately targeted towards tightly defined customer groups. These campaigns can in turn be personalised for delivery through any

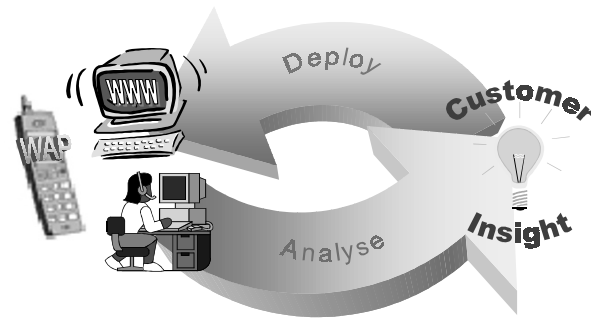


Figure 2 Towards real-time marketing

communication channel, including e-mail, Web and mailshots as well as both outbound and inbound contacts at call centres

- content personalisation: this involves specifying rules that determine how different Web, WAP or interactive digital television users are treated on the organisation's Internet sites, according to their preferences and the organisation's understanding of their profiles and behaviour
- customer interaction management: increasingly, front-office tools enable customers to be handled differently according to what the business knows about them — eg a company might want to treat high-value customers or those at risk of defecting to another supplier, with greater priority than other customers.

While the latter two applications may fall under the banner of service or sales functions, it is a strategic marketing responsibility to determine the way that customer insight (eg value segmentation, defection warnings) is deployed through these technologies. This is in order to personalise the customer experience and drive up customer value. Thus when campaigns are referred to in this paper, this includes personalisation initiatives as

well as conventional marketing campaigns.

And what about the future? The answer is real-time marketing. This involves capturing customer data, analysing them in real-time, generating insight and deploying them 'on-the-fly'. This will allow businesses to use what they learn about the customer during an interaction to affect the outcome of that interaction, not just in contacts that occur days or weeks afterwards. The technology to achieve real-time marketing exists, but the methodologies and discipline required to apply it effectively are still immature.

WHY CLOSE THE LOOP?

It is worth first considering what the current situation is in a typical organisation with large numbers of customers. Many such organisations have data warehouses and analytic capabilities, and use the types of marketing tools described above. But the way that these technologies are used together is typically unsystematic, immeasurable and often ineffective. Here are some common examples of problems with the way that these tools are used:

- customer segmentation strategies are defined at a high level but not

- applied consistently to customer interactions, eg high-value customers are treated the same as low-value customers
- campaign target lists are generated 'ad hoc', manually loaded into marketing databases, and rarely refreshed, target models are redeveloped even less frequently, if ever
- conflicts occur over customer and channel ownership (corporate marketing manager to product marketing manager: 'They're my customers! You can't send them your mailshots!')
- direct responses to campaigns can be tracked, but not indirect responses (ie changes of behaviour)
- campaign performance is measured by response uplift, not return on investment (ie the increase of customer lifetime value attributed to the campaign).

Businesses are able to close the CRM loop by developing processes and infrastructure that can deploy customer insight directly into marketing tools. Furthermore, this approach supports continuous measurement of campaign effectiveness, and iterative campaign targeting improvement. The methodology covers both business processes and technical architecture. As a result, the following benefits can be attained:

- full visibility of campaign performance and return on investment, using both direct and indirect response data, and supporting measurement against actual customer value over time
- enhanced campaign performance, through incremental improvements to target accuracy
- reduced time to market and costs for each campaign, by integrating the

- business and technical architecture necessary to support the campaign management process
- optimised channel utilisation through clear definitions of business ownership and effective management of conflicts.

BUSINESS ARCHITECTURE

To deliver a closed-loop CRM capability, a fundamental culture change must take place in the way that campaigns are run. To achieve this, a number of key preparatory steps are required, as follows:

- define clear ownership of key business 'entities' customers (eg corporate marketing manager, consumer marketing manager); marketing initiatives (eg product manager, corporate marketing manager, sales manager, loyalty programme manager); products (eg product manager); channels (eg marketing communications manager, call centre manager, website manager)
- ensure the right supporting teams are in place: analytics; campaign/content management
- formalise the campaign management process.

The last step involves developing a suitable business process that covers the full campaign lifecycle, from the initial idea to live operation and beyond. An example that can be used as a 'template' for campaign management is illustrated in Figure 3. Figure 3 shows how each activity within the process is clearly defined in terms of the team members who must be involved, the input required, the questions that must be asked during the activity, and the activity's output. Setting out this process in detail is the only way to ensure that the potential benefits of the technologies involved will be realised.

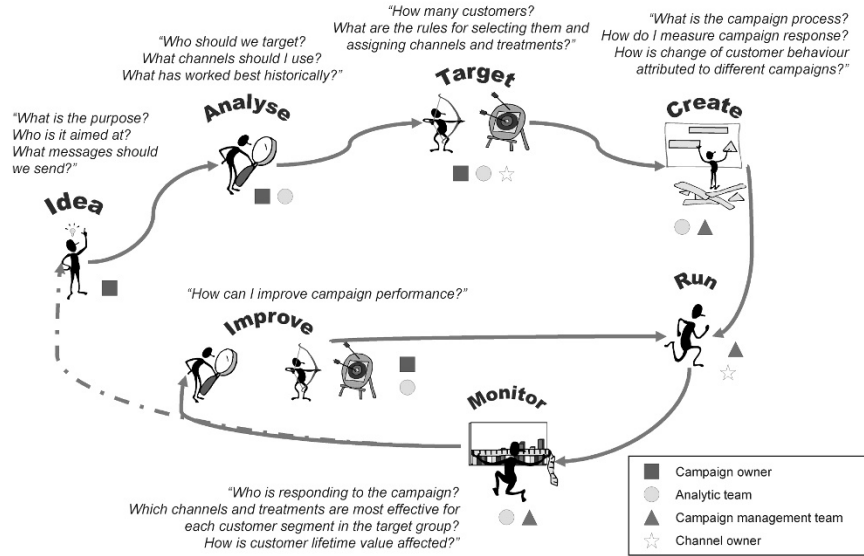


Figure 3 Template campaign management process

TECHNICAL ARCHITECTURE

Figure 4 illustrates a technical architecture to support closed-loop CRM.

At its heart, customer insight is enabled by a historical data repository that supports a range of analytical purposes and techniques. This is referred to variously as a 'data mart' or a 'data warehouse' depending on the level of complexity of data loaded and maintained. Through a series of ETL (extract, transform, load) processes, data feeds from operational systems are transformed into coherent, time-oriented stores of data that can be summarised in different ways to support various analytical applications. The latter may include straightforward management information reporting, on-line analytical processing (OLAP), statistical analysis and data mining.

To support effective operational CRM, the development of an 'operational customer view' (OCV) is necessary to support customer interactions. This may take the form of a simple database of high-level customer data with 'foreign

key' references to detailed data held in other operational systems such as billing, enterprise resource planning (SAP, Baan, etc), and other 'legacy' applications. In more complex environments it may be implemented using Enterprise Application Integration (EAI) technology — providing a 'middleware' layer to access these systems through a single technical interface.

The purpose of the OCV is to provide a single consistent view of the customer that can be used by all customer contact systems to support customer interactions. As well as basic customer data, the OCV should hold data on contact history, campaigns, campaign targets (ie which customers have been sent campaign messages), campaign responses, and data protection flags (eg opt outs).

Marketing tools also need their own database to operate from, referred to here as the 'marketing mart'. This holds data about customers, campaigns, targets and responses. In addition to drawing data from the marketing mart, it is often appropriate to integrate the marketing

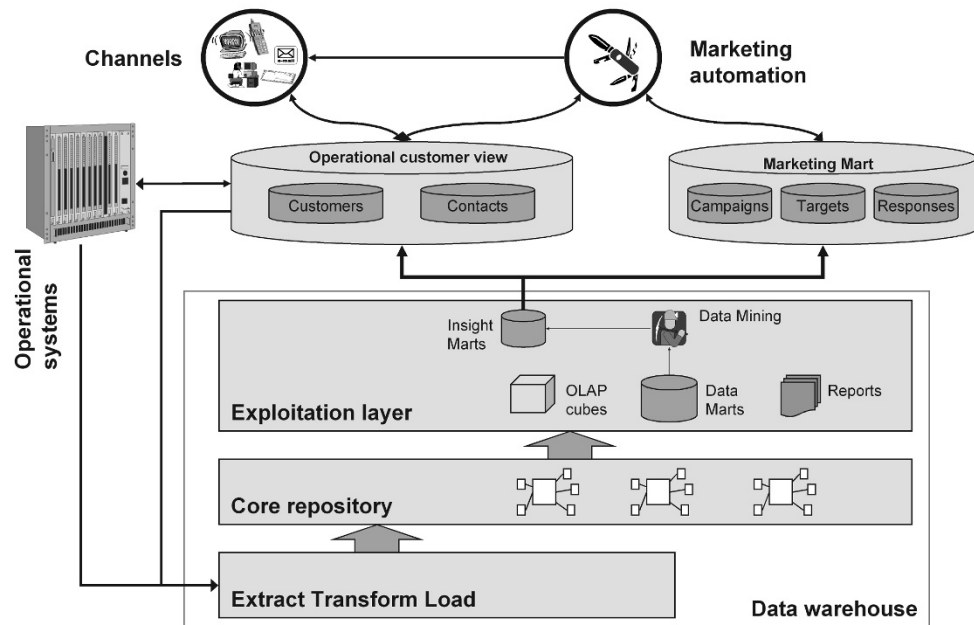


Figure 4 Closed loop CRM technical architecture

automation tools directly with the OCV where possible to ensure that:

- customer data used by marketing tools is consistent and up-to-date with other systems, eg ensuring that changes of address or ‘do not contact’ opt-out flags are applied immediately for each campaign execution rather than waiting for a refresh from the data warehouse
- campaign action and response data can be transferred seamlessly between the marketing tools and customer interaction systems.

Data from the OCV must be loaded regularly into the data warehouse core repository using ETL tools. Within the data warehouse, visibility and analysis of campaigns should be supported with a historical store of campaign targeting data. This will show how customers have been targeted through particular campaigns, treatments and channels, and indicate how each customer responded

(if at all). Each record may also hold a ‘score’ representing the predicted probability that the customer will respond to the campaign when delivering that particular treatment over the selected channel. This is illustrated in Figure 5 as a ‘star schema’ for linking each of the various dimensions relating to campaign management.

Customer insight generated through data analysis activities is normally held in data marts within the data warehouse. These data must be transferred to the OCV and/or marketing mart as appropriate — this is explained in more detail in the section on targeting below. Some examples of the types of insight applicable are as follows:

- strategic segmentation, such as sociodemographic, geodemographic and value based
- loyalty indicators, how likely a customer is to defect to another supplier
- channel propensity, how likely a

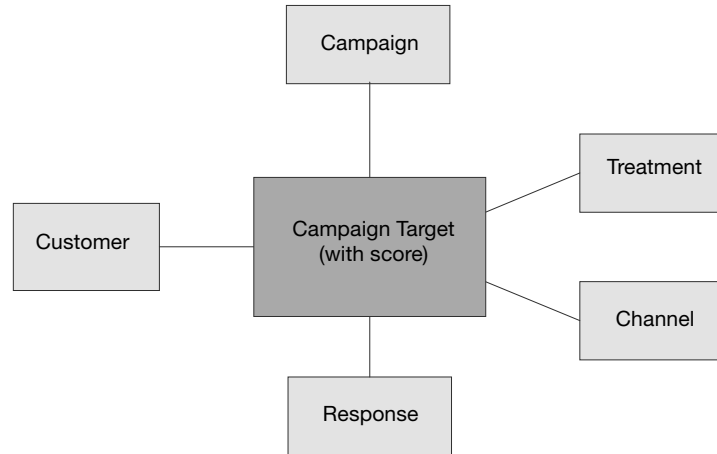


Figure 5 Campaign targeting data in the data warehouse

- customer is to use particular channels such as the Web, email, etc
- campaign propensity scores, how likely a customer is to respond to a particular campaign
- response value scores, the estimated value of a positive response to a campaign.

A RECIPE FOR SUCCESS

There are a number of points that need to be considered extremely carefully at the outset of any customer insight project. The key elements are:

- selecting targets for campaigns
- tracking responses
- measuring performance.

These factors are now discussed in more detail.

TARGETING

What is a campaign target? As explained previously, this is a combination of:

- a campaign
- a customer
- a channel

- a ‘treatment’, the message to be delivered.

To ensure that channels and treatments are chosen appropriately and utilised effectively, it is crucial to ensure that:

- historical channel effectiveness is analysed by customer segment and campaign type to feed into decisions on channel selection for future campaigns
- the campaign management team establishes an SLA with each channel owner. For example, this may require a print bureau to deliver 100,000 mailshots per week, or a call centre to deliver 10,000 outbound marketing calls per week
- channel owners provide visibility of channel utilisation to the campaign management team
- restrictions on how often customers can be contacted through each outbound channel are agreed with customer owners
- the campaign management team takes on responsibility for optimising the use of every channel and customer contact ‘window of opportunity’, based on channel availability, the

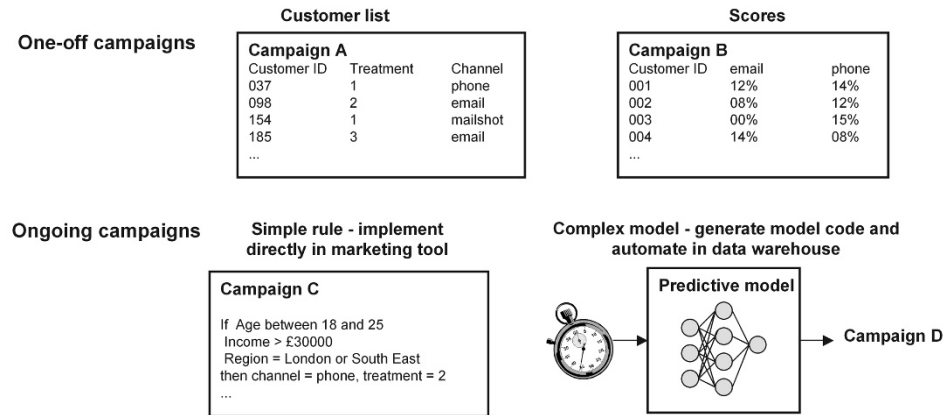


Figure 6 Methods for specifying target groups

demand placed on each channel by different campaigns, and constraints applicable to customer contact frequency.

Selection of treatments is likely to be based on intuition at the initial stage of a campaign. Provided, however, that the right data are made available and analysed effectively, selection of treatments can be optimised during ongoing campaigns.

Of course, disagreements can and do arise over what is the best campaign to deliver to a customer, or how a particular channel should be used. Some campaign management tools provide the option to choose customer targets automatically in order to optimise revenue or growth given a combination of propensity scores and constraints such as channel capacity and maximum frequency of customer contact. It requires, however, a cultural leap of faith to entrust the technology with marketing decisions at this level. In practice, it is imperative to ensure that a rapid and effective escalation route exists to resolve these kinds of conflict.

Finally, the use of control groups is crucial to ensure that campaign performance can be measured and

improved. These come in two forms:

- inclusive: additional randomly selected targets added to the target group, to evaluate the ‘uplift’ of the targeted selection versus an untargeted selection. This allows the benefit of the propensity modelling process to be measured
- exclusive: targets randomly removed from the target group, to ensure that the effect of the campaign on customer value can be measured by comparing the ongoing value of customers who were targeted versus those who were not.

Given a model that can be used to generate the target and control groups, the question of how to transfer this information to the campaign management tool remains. There are a number of methods, as illustrated in Figure 6 and Figure 7. The approach taken will depend on whether this is a one-off campaign or an ongoing campaign, and possibly on the specific marketing automation technology being used.

For a one-off campaign, the most straightforward approach is to export a list of targets and load these directly into

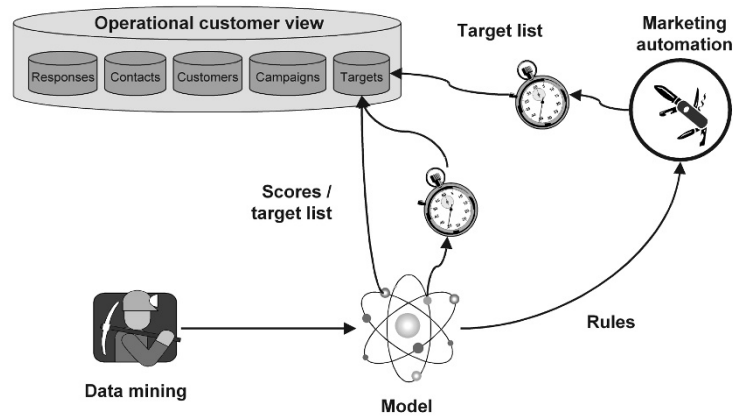


Figure 7 Approaches for loading targets into the marketing tool target database

the marketing mart. The exported data may be either a simple list of targets, or a list of potential targets with scores. In the latter case scores may be used by the marketing tool to decide which customers to target, the basis of estimated cost and revenue data managed.

For ongoing campaigns, businesses can benefit from automating this process. There are two possible approaches — with the option selected depending on the complexity of the model and the capabilities of the marketing tool:

- for simple models (eg ‘all customers over 50 years old without credit cards’), it may be possible to implement the rule in the marketing tool to enable it to select the targets at the time of campaign execution
- for more complex models (eg a neural net or other ‘black box’ model), it will be necessary to automate the analytical routines that generate the targets from the data and export these to the marketing tool.

The last option presents two further possibilities — ‘off-line scoring’ or ‘just-in-time scoring’. Off-line scoring involves setting up a batch scheduler to run the analytic model that generates the scores regularly (eg once per week) as

required, and then to push the data into the marketing mart. While this is invariably possible from a technical viewpoint, it requires involvement from the data warehouse administrator to schedule the job.

Conversely, the ‘just-in-time’ approach involves selecting the model to apply in the marketing tool itself. The tool then connects to the analytical software and runs the scoring routine immediately before campaign execution, or at the point of user interaction for real-time applications. This has the advantage that the scores are always generated using the most up-to-date data. A few software toolsets provide this sort of integration ‘out-of-the-box’. Provided that companies buy the analytical software and marketing software from the same vendor, these tools provide a simple user interface to select models to apply with no need for IT involvement.

The routes through which the different approaches map onto the technical architecture are illustrated in Figure 7.

RESPONSE TRACKING

What is a response to a campaign? In some cases, the answer is less than

obvious. In general responses fall into two categories:

- direct response (where the customer explicitly responds to a campaign)
- indirect response (where the response takes the form of a change of behaviour, eg the purchase of a product, or the use of a new service).

The problems to be addressed are:

- the measurement of indirect responses
- how to attribute responses between multiple campaigns with similar goals.

The first problem is particularly relevant to mass media campaigns that aim to deliver long-term value through increased brand awareness. It is necessary to define the business rules for identifying responses to these campaigns at the outset, and to apply these rules to evaluate returns in the long term.

The second problem is more difficult in situations where multiple campaigns can elicit similar indirect responses. For example, one campaign may encourage people to buy a new product, whereas another campaign may encourage people to visit a particular shop. In this situation it is vital to seek agreement between campaign owners on how to attribute responses to campaigns before the campaigns are implemented.

MEASURING PERFORMANCE

To measure campaign performance, it is important to draw a distinction between the short and long-term goals of a campaign:

- in the short term, a campaign should elicit responses, a high performing campaign is one which generates a high response rate

- in the longer term, a campaign should increase customer value, a high performing campaign substantially increases the overall profit generated by the targeted customers over their lifetime.

In some cases these goals are the same. For example, if a campaign's aim is just to encourage customers to buy a product as a 'one-off' event, the boost to the customer's lifetime value is just the profit generated by that sale.

In other cases, the estimate of the value generated by a positive response is frequently inaccurate or misleading. For example, if a campaign offers a free trial period, will the customers responding to that campaign be as profitable in the long term as the average customer?

This is addressed by establishing a customer value model, which aims to predict each customer's overall value to the organisation over the lifetime of the customer. This is often the subject of a separate and potentially complex exercise, but gaining this customer understanding is one of the key drivers behind building a customer insight capability in the first place. Put simply, the customer value model is a formula for estimating the future revenue or profit to be derived from each individual customer, expressed in terms of a set of factors such as:

- how long the customer is expected to remain loyal to the organisation
- which product types they currently purchase, and how much profit/revenue is generated by each such product
- what the likelihood is that they will purchase each additional product type, and what the predicted profit/revenue to be generated by each such product is
- what the predicted future cost is of servicing, retaining and developing

the customer, including the cost of service, sales and marketing.

The model needs to be defined as an algorithm applied to available customer data. Once defined, it can be implemented in the data warehouse to re-calculate the value of each customer at regular intervals. By combining the customer value data with campaign targeting and response data for target and control groups, businesses will be ready to evaluate campaign performance in terms of its effects on customer value. Of course, to get meaningful results this can require long-term monitoring.

CONCLUSION

Many organisations have tried marketing automation technologies, but few implementations deliver the rewards promised. The most common reason for this is the difficulty that organisations face in managing the substantial cultural shift necessary to make effective use of the technology.

In the cases where the technology delivers real benefit, the tools are just elements within a well-managed human process, one in which ownership and responsibility for controlling campaigns, channels and communication are well defined, understood and put into practice.

The tools themselves are gradually maturing, but there is still a clear divide between the offline and online marketing tools. Campaign management tools focus on the campaign process as an offline

activity, whereas content customisation engines emphasise support for online interactions over electronic channels. The convergence of the two types of tool will enable users to design campaigns targeted across all channels, tracking offline channel responses while reacting to online channel interactions in real-time. A few emerging technologies claim to inhabit both worlds, but the best-of-breed tools are each firmly camped in one or the other.

As for closed-loop real-time marketing, the benefits have still to be proven. While the concept may appear attractive in the go-faster world of e-commerce, it is important to remember that customer insight is a human process that delivers most of its value by letting businesses know their customers better. By taking the human brain out of the loop, real-time marketing loses out on the most powerful analytic engine on the planet.

CRM does not, however, begin and end with marketing. The approach described in this paper explains just one example of a process to be implemented within a broader CRM strategy. Ultimately, the most rewarding investment to be made in CRM is in delivering change in a way that alters the attitude of the entire organisation to become 'customer-centric'. By providing visibility of the customer across the entire enterprise, every interaction with the customer can be coordinated, providing the key to increased customer value and profitability.

©Detica Ltd