

Cadence : Robotic Data Scientist for Marketeers by The Data Team



Efficiency of marketing campaigns leading to customer delight is a great challenge for B2C marketers at BFSI companies



Challenges

- Not knowing the customer well enough
- Contextually unaware decisions
- Lag in insight calculation and implementation
- Inability to scale to complete customer base

Majority of consumers say personalization plays a role in their purchasing decisions



Ideal Solution

Dynamic profiling of customers which include socio-economic and behavioral data
Advances self-learning domain aware models

Three-fourth of consumers prefer to do business with brands that use personal information



Desired Outcomes

- Contextual and timely decisions
- Personalized and enhanced customer experience
- Scalable data models enabling dynamic processing that provide instant domain-aware decisions

Enhanced customer satisfaction and reduced time to market



Cadence by The Data Team



Get holistic view of customer and automatically target each customer with the most relevant product or service at the precise time

Know your Customers

- Holistic view of customer: Socio-economic , behavioral
- Micro-segments and individual – level characteristics

Contextual decisions

- Domain-aware dynamic decisions powered by ontologies
- Merging digital and physical contexts

Timely Insight and Implementation

- Intelligent audience discovery and roll out real time offers
- Integration with existing channels and campaign management systems

Provide personalized services or offers by targeting customers based on micro-segments and individual level characteristics

Cadence with Microsoft Azure



Scalable, real time and contextual marketing engine to enable personalized service offerings. Elasticity of data storage fostering dynamic remodeling of data silos. Execute marketing campaigns through existing channels and systems

Solution Alignment

Cadence with Microsoft Azure

Powered by Azure big data stack, solution allows scalability of complete customer base



Cadence with Microsoft Azure

Operationalize marketing analytics through domain aware self-learning models enabled by machine learning on Azure



Cadence with Microsoft Azure

Combine data from both external and internal data sources ; integrate with existing channels and systems



Case Study : Real Time Offers



A global MNC bank in Singapore was looking for a suitable solution to increase usage of its credit cards by its customers. Non-scalable technology and coarse segmented targeting are the major challenges.

Vertical : Banking

Region : Singapore

The retail banking division wanted to send out offers in real-time to eligible customers based on their propensity to use up the offer. The business benefit from rolling these offers out is to encourage customers to spend more on credit cards.

Cadence, the solution provided to marketing teams allowed them to target their customers instantaneously with real-time offers rolling out that are contextual and location-aware. This cloud based solution allows scalability and helps in operationalizing marketing analytics through domain-aware self-learning models.

Win Results

Marketing teams were able to fetch an holistic view of customers and is used to execute marketing campaigns dynamically and contextually

Personalized, real-time, contextual and location-based offers are now made available to huge volumes of credit card customers instantaneously.

Cadence served as an efficient conversion tool which increased credit cards usage and thus enhanced RoI. Rewarded as one of the innovative marketing campaigns in the region.

BOM



Azure consumption will need to be computed in detail but just the Cadence solution alone has the potential to drive about \$50k per year for a mid-tier bank. With all the solution examples we have put together, the potential is easily north of \$100k per bank.

Processing Type	Service type	Description
Batch-based	Storage	XXX TB
	Data Factory	NN hours/month of data movement from on-premise
	HDInsight for batch analysis	NN worker nodes, 8 cores each (total 96 cores) running 8 hours/day to process data in batches
	Machine Learning (training)	Leveraging SparkML included in HDInsight to train and persist the models
Real-time	Machine Learning (scoring)	Leveraging SparkML included in HDInsight to load the persisted models and score the new data points
	Azure Cosmos DB	MM GB SSD, Request units AAA k/sec
	Event Hubs	hundreds of million events/month
	HDInsight master node	NN worker nodes, 16 cores each (total 32 cores) running 24/7
	HDInsight for real-time analysis	Leverage Spark Streaming included in HDInsight Average XXX cores consumption per day
	Virtual Machines (to configure business rules)	
	VPN Gateway	High performance VPN gateway, sending outbound messages to fulfillment for a total volume of XXX GB/month