Teradata – Microsoft® Azure Reference IoT Architecture

Artur Borycki, Teradata Technology & Innovation Office

11.17 EB7303 INTERNET OF THINGS

Microsoft TERADATA.

Table of Contents

Meeting IoT Needs of the Organization

Organizations are understandably excited about the immense value to be derived from harnessing the raw data gathered from the Internet of Things (IoT). IoT data can produce high-value insights, especially when processed with real-time and near real-time analytics. Whether the IoT devices are reporting from industrial processes, biometrics from medical devices, consumer behavior, or a spectrum of other realms, big data from IoT—whether analyzed on its own or combined with other sources or applications—can yield high-value metrics and insights.

Of course this creates a challenge for IT—translating into reality the vision of becoming a data-driven organization. IDC estimates that "80 percent of companies suffer from unclear IoT goals, narrow use cases, and limited analytics scope."

Yet the promise of IoT is so great that Gartner foresees a compound annual growth rate of 20 percent resulting in $7.1 trillion IoT market by 2020.

The good news for IT and others tasked with making efficient IoT a reality is that Teradata and Microsoft Azure provide an ideal platform for cloud-based IoT solutions.

Value Across the Organization

IoT, when robustly deployed and fully-integrated across multiple data sources—including traditional relational databases and unstructured and semi-structured data stores—and powered by real-time and near real-time analytics, changes everything.

Some basic examples of improvements due to IoT analytical deployments include:

- **Products and Services**. Periodic plans and designs can be evolved to continuous design and feedback based on IoT analytics.
- **Operations**. Silos and manual operations can be replaced with automation and feedback loops.
- **Support**. Unexpected failures can become less common as IoT analytics allow processes and products to be fixed before something breaks.
- **Business Goals**. Traditional product-centered planning can be transformed, generating new digital business revenue.

All of these opportunities—and benefits—are dependent, though, on deploying a robust and tightly integrated architecture for harvesting and analyzing data from an array of devices and the ability to seamlessly integrate with a range of data stores, services, and applications.

Better Together: Teradata and Azure

Teradata analytics on Microsoft Azure provides an ideal cloud-based deployment scenario for organizations seeking to harvest the benefits of IoT and big data. And Teradata on Azure also provides the complete integration and robustness that IT requires when they are asked by the business side of the organization to "make IoT happen."

TERADATA.

Teradata – Azure Reference Architecture

Azure provides the most complete platform for cloud-based solutions, while Teradata® Database and Teradata Aster® Analytics, combined through the Teradata Unified Data Architecture™, have proven themselves with some of the world's largest data management deployments.

When considering IoT architecture, three logical segments are:

- Device connectivity and management.
- Data processing, analytics, and management.
- Business connectivity.

While touching briefly on the other areas, this paper focuses on the data processing, analytics, and management elements of the Azure IoT Reference Architecture.

Device Connectivity and Management

As shown in Figure 1 (callout numbers 1 through 4), Microsoft Azure provides a solid front-end to the IoT reference architecture, with an Azure-powered cloud gateway connecting IoT clients and other IP-capable devices. Azure also supports connectivity with custom gateway protocols.

Azure IoT Gateway

The Azure IoT protocol gateway provides a secure, end-to-end framework for protocol adaptation, designed for high-scale, bi-directional device communication with the Azure IoT Hub. The Azure IoT protocol gateway includes a Message Queue Telemetry Transport (MQTT) protocol adapter that enables you to customize the MQTT protocol behavior if necessary.

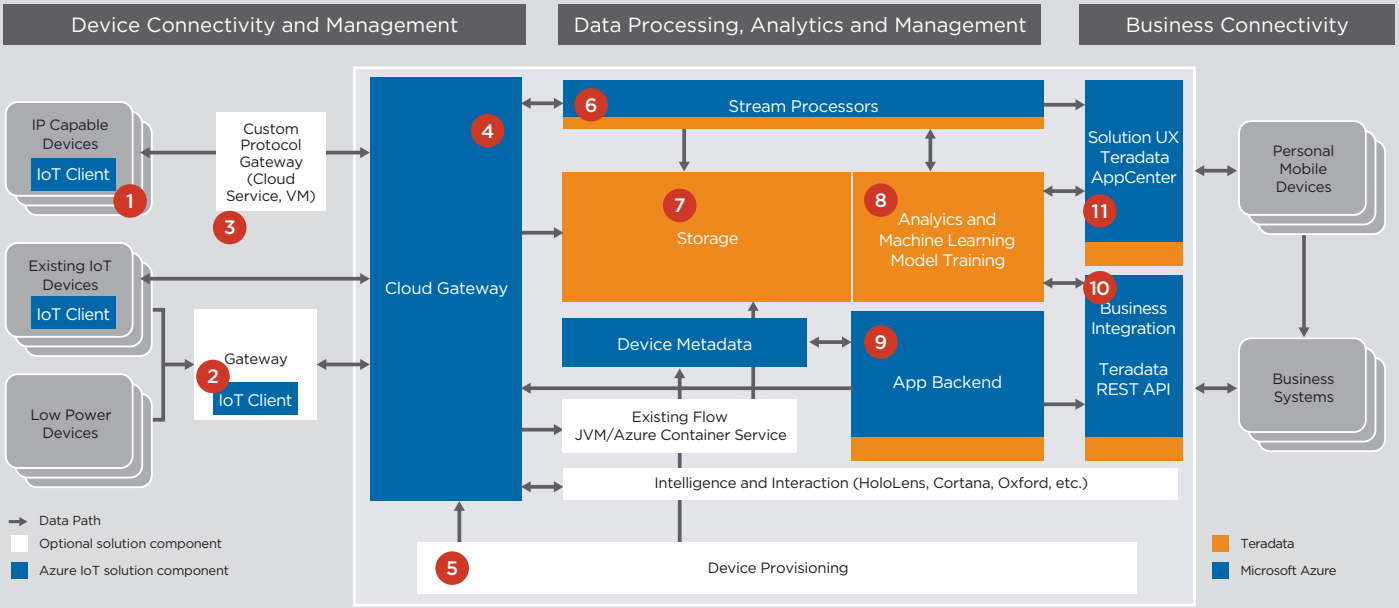

Figure 1. Azure – Teradata IoT Reference Architecture.

Azure IoT Hub

Azure IoT Hub is a fully-managed service that enables reliable and secure bi-directional communications between millions of IoT devices and a solution back end. Azure IoT Hub:

- Provides multiple device-to-cloud and cloud-to-device communication options, including one-way messaging, file transfer, and request-reply methods.
- Provides built-in declarative message routing to other Azure services.
- Provides a queryable store for device metadata and synchronized state information.
- Enables secure communications and access control using per-device security keys or X.509 certificates.
- Provides extensive monitoring for device connectivity and device identity management events.
- Includes device libraries for the most popular languages and platforms.

Processing, Analytics, and Management

Processing, analytics, and management, as shown in Figure 1 (callout numbers 5 through 11), are at the heart of the Azure – Teradata reference architecture. Basic architectural elements include:

- Data ingestion
- Stream processors
- Storage
- Analytics, machine learning, and model training
- Application backend
- Solution UX
- Business integration

Data Ingestion with Teradata Listener™

An early architectural consideration is the ingestion framework. Teradata Listener™ is an ingestion framework that takes data from the Azure IoT Hub into the Teradata ecosystem, supporting the protocols and APIs provided by the Azure IoT hub. Listener gives you the ability to move multiple data streams from the Azure gateway to storage—whether on Teradata Database, Apache® Hadoop®, or other data stores.

Multiple Sources and Targets

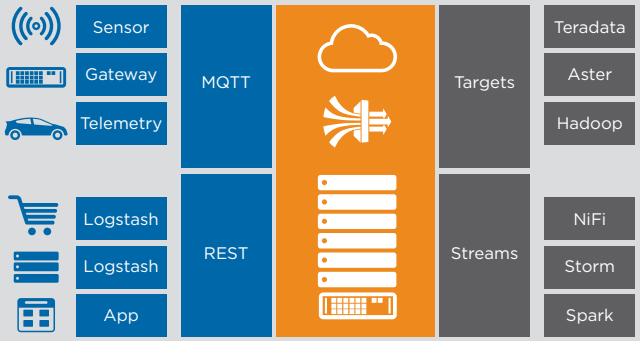

Teradata Listener unifies the big data ingestion process by capturing multiple, high-volume data streams continuously from a variety of sources, and persisting them into one or more of the data stores. Listener is an intelligent, self-service solution for ingesting and distributing extremely fast-moving data streams throughout the analytical ecosystem.

Using Listener to capture data from multiple data streams, both internal and external, enables businesses to listen to their data and take actions more rapidly than ever, leading to better business outcomes.

Self-Service Solution

Teradata Listener provides a self-service software solution for programmers that minimizes the complexity of building and supporting data streams and provides the capabilities to:

- Operate as a central service to ingest and write hundreds of streams.
- Reliably deliver data without loss.
- Provide low-latency ingestion for near real-time applications.
- Run as a software-only solution in the data center or in the cloud.

Teradata Listener's architecture enables the decoupling of incoming data streams with the outgoing distribution processes. Listener buffers the distribution output intelligently when the target systems are full and activates the distribution later when the target system allows for it—all without any manual intervention.

Teradata Listener™ Dashboard

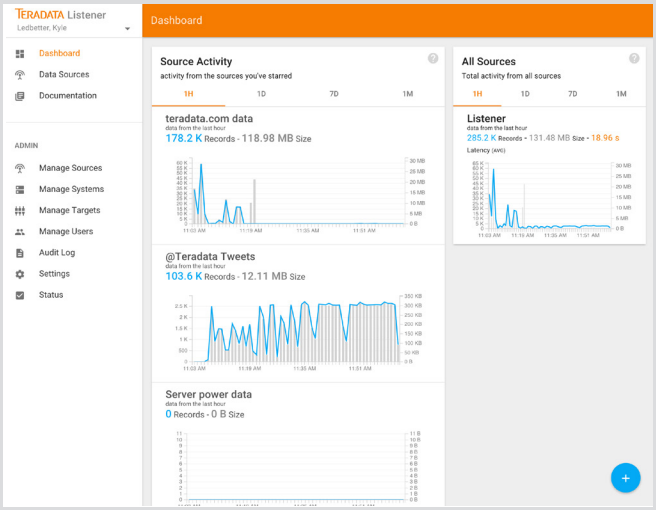

Listener's ingestion services are accessible by popular interfaces, including:

- **Representational State Transfer (REST)**. REST, the popular http transport protocol, is a universally accepted protocol for modern applications.
- **MQTT**. MQTT is the ideal protocol to overcome the emerging mobile-to-mobile (M2M) or IoT world of connected devices where bandwidth and battery power are at a premium. Any user can easily invoke the Listener's ingestion services to send continuous data streams to a data warehouse, discovery platform, or Hadoop platforms.
- **Application Programming Interface (API)**. APIs provide more flexibility to developers to access the data flowing through Listener. This data from Listener can be used to create custom dashboards or integrate with other processing engines for transformations.

Enterprise-Grade Reliability with Intelligent Buffering and Failover

Teradata Listener is built to reliably deliver data without loss. When the target systems are full, and therefore cannot accept any more data distribution, Listener's intelligent buffering mechanism intermittently persists data within the system and delivers the data later when the target becomes available, ensuring no single data stream is lost ever. Listener's inbuilt failover architecture deals with failures automatically. If any service on Listener goes down for any reason, it is automatically re-started on the server, ensuring continuous availability of all the services. Listener maintains redundant copies of the data throughout the data ingestion and distribution process in case the first copy is lost.

Stream Processors

Stream processing—a method of continuous computation of incoming data streams, which can generate results in milliseconds—can pull immediate value from IoT information, including process monitoring and triggering threshold-event responses.

Stream processors perform real-time data processing at scale and support:

- Data aggregation
- Data enrichment
- Complex event processing

Teradata integrates with stream processors including complex event processing (CEP) to help organizations to analyze their data in motion and at rest in an integrated analytics and data management architecture.

The Teradata Aster® Scoring SDK is intended for systems such as stream processors that follow events in real-time and must take action based on these events in real-time with the support of analytics. Teradata Aster Scoring SDK applies predictive analytics to make timely decisions based on real events. Teradata Aster Scoring SDK also makes Teradata Aster Analytics functions available for real-time prediction. Some use cases for Teradata Aster Scoring SDK include:

- Fraud prevention
- Churn reduction
- System failure predictions
- Site personalization
- Purchase recommendations
- Dynamic promotion pricing

Storage

From an architectural standpoint, storage should be considered in terms of integration between existing enterprise data stores, cloud-based resources, and big data deployments such as Hadoop, Apache Cassandra, or Apache HBase™.

Relational Databases

Relational databases continue as the foundational analytics technology for the enterprise. The strong data types and schema support rich querying and reporting, and security can be applied down to the row and cell level. With more than 35 years of cutting-edge development, Teradata Database is the market-leading platform for delivering strategic and operational analytics throughout the organization, providing users with a single source of consistent, centralized, integrated data with petabyte scalability. Teradata Database on Azure combines Teradata's industry-leading technology with the agility and flexibility of Microsoft Azure infrastructure.

Big Data Storage

While Relational Database Management System (RDMS) solutions are powerful tools for structured data, the realm of big data is filled with unstructured and semi-structured data that is more efficiently stored and queried by using Hadoop and similar technologies.

Teradata's Portfolio for Hadoop is a flexible suite of hardware, software, and services for organizations to integrate enterprise Hadoop into a Teradata environment across a broader enterprise architecture. Hadoop continues to gain momentum in the enterprise by providing a cost-effective, massively scalable environment for loading, storing, refining, and exploring very large amounts of data of any format.

Teradata has deep engineering partnerships with Hortonworks®, Cloudera® and MapR offering customers the ability to choose which Hadoop distribution best meets their needs, while ensuring integration capabilities and the advantage of world-class Teradata service and support.

Teradata Appliance for Hadoop is a purpose-built, integrated hardware and software solution for data at scale. The appliance is offered with flexible appliance configurations with dual twelve-core Intel® Xeon® processors at 2.5GHz or dual eight-core Intel Xeon processors at 2.6GHz with configurable memory options from 128-512GB. An optimized version of Hadoop software from Hortonworks and Cloudera, SUSE Linux 11, connectors for high-speed data transfer, and a 40Gb/s InfiniBand Teradata BYNET® V5 network all combine to deliver the fastest time to data access.

Teradata Appliance for Hadoop is linearly scalable to petabytes of data to meet big data storage requirements. Each node is equipped with 4TB capacity HDD drives for your data storage needs.

By minimizing the number of moving parts required for deployment and operations, the appliance allows companies to achieve faster time to value by just plugging the appliance into existing infrastructure, thereby leveraging current investments in technology and resources.

Key-Value Databases

Key-value databases, such as Cassandra and Apache HBase, are efficient for storing collections of semi-structured objects and records which can have many fields. Key-value databases provide flexibility, scalability, and high performance that make them attractive for cloud-based storage. Teradata has deep experience building solutions with the open source technologies such as Cassandra and HBase, which can play a role in overall IoT architecture.

Document Databases

Document databases, typically using the JSON open-standard file format, allow storage of—and querying against—text-based files such as documents or website content. Document databases help widen the data lake that organizations can draw from, though there are limitations on joins and other advanced operations. Teradata supports JavaScript Object Notation (JSON) as a native data type, which allows storage of documents as JSON objects within Teradata, making the document information available for combining with relational for analytics. Azure Cosmos DB, Microsoft's globally distributed, multi-model database, natively supports multiple data models, including but not limited to document, graph, key-value, table, and columnar data models.

Analytics, Machine Learning, and Model Training

The first task is to collect data; the second is to derive value from it. This is where analytics—including those applied to machine learning and model training—come into play.

Analytics can be roughly divided between real-time and at-rest. Real-time and near real-time analytics—sometimes called hot analytics—are performed while data is still in transport, flowing in data streams from source end points or gateways to storage.

Real-time analytics can be critically important because some data events have very perishable value with respect to time, whereby failure to take action within milliseconds makes all the difference. Examples include a request for payment that may be fraudulent, or sensor reading that may indicate a pending mechanical failure.

Real-time capabilities often benefit from tight integration with historical data and analytical tools to improve context and rapidly iterate on enhancing real time algorithms.

Real-Time Analytics with Teradata QueryGrid™ and Presto

When data is persisted in multiple engines, a virtual query capability is needed to enable users and applications to connect-the-dots while minimizing data movement. Teradata QueryGrid™ and Teradata's use of Presto lets organizations work with a seamless data fabric across all of their data and analytical engines for no-hassle analytics.

Presto, first developed by Facebook, can, with a single query, combine data from multiple sources. Presto offers connectors to data sources including files in a Hadoop distributed file system, Microsoft SQL Server, Azure SQL Database, Azure Cosmos DB, MySQL, Apache Kafka, Cassandra, PostgreSQL and Redis.

QueryGrid enables a data fabric with seamless multi-system analytics, transparently harnessing the combined power of multiple analytic engines to address a business question. QueryGrid enables high performance query plans using data from other sources while using systems within the Teradata Unified Data Architecture™ such as passing workload priorities to achieve the best use of available resources. QueryGrid, for example, makes it possible to launch on-the-fly querying of data in Hadoop.

From an IT perspective, QueryGrid offers benefits including the ability to:

- Automate and optimize use of your analytic systems through "push-down" processing across platforms.
- Minimize data movement and process data where it resides.
- Minimize data duplication.

Teradata Aster Analytic Functions

STATISTICS
- AdaBoost
- Approximate Distinct Count
- Approximate Percentile
- CMAVG
- ConfusionMatrix
- ConfusionMatrixPlot
- Correlation
- CoxPH
- CoxPredict
- CoxSurvFit
- Cross Validation
- Distribution Matching
- EMAVG
- Enhanced Histogram
- FMeasure
- GLM
- GLMPredict
- Hidden Markov Model
- Histogram
- KNN
- LARS Functions
- LinReg
- Logistic Predict*
- Logistic Regression*
- LRTEST
- Non-linear Kernel SVM
- Percentile
- Principal Component Analysis
- Random Sample
- Sample
- Shapley Value
- SMAVG
- Support Vector Machines
- VectorDistance
- VWAP
- WMAVG

CLUSTER
- Canopy
- Gaussian Mixture Model
- Kmeans
- KMeansPlot
- Categorical Clustering
- Minhash

DATA TRANSFORMATION
- Antiselect
- Apache Log Parser
- FellegiSunterTrainer
- FellegiSunterPredict
- IdentityMatch
- IpGeo
- JSONParser
- Multicase
- MurmurHash
- Number as Categories
- OutlierFilter
- Pack
- Pivot
- PSTParserAFS
- Scale Functions
- StringSimilarity
- Unpack
- Unpivot
- URIPack
- URIUnpack
- XMLParser
- XMLRelation

PATH, PATTERN AND TIME SERIES
- Arima
- ArimaPredictor
- Attribution
- Burst
- ChangePoint-Detection
- Causality Detection
- DTW
- DWT
- DWT2D
- FrequentPaths
- IDWT
- IDWT2D
- Interpolator
- nPath
- Path_Analyzer
- Path_Generator
- Path_Start
- Path_Summarizer
- SAX
- SAX2
- SeriesSplitter
- Sessionization
- Shapelets
- Unsupervised Shapelets
- Varmax

VISUALIZATION
- CfilterViz
- NpathViz

*depreciated

ASSOCIATION
- Basket_Generator
- Cfilter
- FPGrowth
- KNN Recommender
- WSRecommender

DECISION TREE
- Forest_Drive
- Forest_Predict
- Forest_Analyze
- Single_Tree_Drive
- Single_Tree_Predict

NAÏVE BAYES
- naiveBayesMap
- naiveBayesReduce
- naiveBayesPredict

DEEP LEARNING
- Neural Networks

TEXT
- Chinese Text Segmentation
- LDA Functions
- Levenshtein Distance
- Named Entity Recognition (CRF Model)
- Named Entity Recognition (Max Entropy Model)
- nGram
- PoSTagger
- Sentenizer
- Sentiment Extraction Functions
- Text Classifier
- Text_Parser
- TextChunker
- TextMorph
- TextTagging
- TextTokenizer
- TF_IDF

TERADATA ASTER DATABASE SYSTEM FUNCTIONS
- nc_skew
- nc_relationstats

GRAPH
- AllPairsShortestPath
- Betweenness
- Closeness
- EigenvectorCentrality
- gTree
- LocalClustering-Coefficient
- LoopyBelief-Propagation
- Modularity
- nTree
- PageRank
- PersonalizedSALSA
- RandomWalkSample

LOCATION ANALYSIS
- LoadGeometry
- PointinPolygon
- GeometryOverlay

See the Analytics Foundation Guide for a Complete Review of Analytic Operations
info.Teradata.com

- Transparently automate analytic processing and data movement between systems.
- Enable easy bi-directional data movement.
- Integrated processing without administrative challenges.
- Leverage the analytic power and value of your Teradata Database, Teradata Aster Database, open-source Presto and Hive for Hadoop, Oracle Database, and powerful languages such as SAS, Perl, Python, Ruby, and R.

Multi-Genre Advanced Analytics™ with Teradata Aster Analytics

Conventional analytics are those that have been tried and true for decades, such as regression, CHAID, decision trees, and countless others. Emerging analytics stem from big data innovation to enable procedural analytical processing such as graph and text analysis. Multi-genre allows for the integration of all types of analytics into a single processing environment.

Teradata Aster Analytics delivers such Multi-Genre Advanced Analytics™ at scale to empower business users to uncover and operationalize non-intuitive insights. With Aster Analytics, analysts are able to seamlessly combine different analytic techniques against multi-structured data to address any use case within a single solution. This solution includes over 100 prebuilt advanced analytics functions covering multiple analytic genres such as path, pattern, statistics, text, graph and machine learning. It also has business user-friendly interfaces, such as SQL, R and point-and-click applications, that allow analysts to quickly explore, discover, and model data across any source then operationalize business analytics.

From a big data perspective, Aster Analytics can be used for high-value insights, including:

- **Path Analytics**. Path analytics capture the customer's journey to conversion.
- **Text Analytics**. Text analytics show customers sentiment analysis.
- **Graph Analytics**. Graph analytics provide influencer analysis.
- **Machine Learning**. Machine learning identifies patterns that can be used to generate proactive predictions such as to prevent part or process failures.
- **IoT Analytics**. IoT analytics of sensor and machine data.

Application Backend

While much of big data discussions center on accumulation of data and running analytics against it, another important element from an architectural standpoint is having the ability to feed the data into customer relationship management, enterprise resource planning, supply chain management, and other line of business applications.

In addition to Teradata APIs and integrations, organizations can make use of Azure API Management for exposing and managing APIs, or create orchestrations with Azure Logic Apps.

Solution UX

Teradata AppCenter provides an application framework—including the Teradata Analytics Framework—that fosters collaboration and leads to faster time to value. Teradata AppCenter provides a user-friendly web-based environment for the development and deployment of applications across the business community.

The Teradata Analytics Framework enables client-led connectivity, using HTTP to access different technologies and solutions. This facilitates creating applications that connect with Teradata without requiring additional code or client software inside your application.

AppCenter enables quick definition and integration of commonly used analytics and provides a marketplace for reuse and collaborative self-service. Pre-built big data apps are available to solve specific business challenges in industries like consumer financial, entertainment and gaming, healthcare, manufacturing, retail, communications and cable, travel and hospitality.

Teradata also integrates with Microsoft Power BI, a suite of business analytics tools that deliver insights throughout your organization.

Business Integration

Teradata supports business integration of IoT data through a number of pathways, including:

- Teradata Platform Services
- Teradata REST Services
- Teradata Customer Interaction Manager
- Teradata Real-Time Interaction Manager
- Teradata Integration with Microsoft Power BI

Teradata Platform Services

Developers, administrators and analysts can easily create robust data-focused applications using Teradata Platform Services. Platform Services allows for seamless integration with internal and external ecosystem tools, and fosters collaboration and innovation with the Teradata development community.

Teradata's Platform Services are developed with a cloud-first mentality, leveraging micro-services and multiple levels of abstraction to ensure scale across the most diverse set of workloads. Examples of commonly-used services would be for transforming, caching, logging, authenticating, provisioning, scheduling, and monitoring.

Teradata REST Services

Teradata REST Services provides application developers a simplified, modern interface to connect to data from a web page or application. To accelerate the development of external applications using a common interface mechanism to address the Teradata Database, Teradata REST Services is a foundation for supporting Restful APIs into various Teradata products.

Teradata REST Services, in combination with the REST API for Teradata Database, provides the following:

- Allows SQL to be submitted to Teradata via a RESTful web service.
- Uses JSON as the data-interchange format and HTTP as the protocol.
- Provides a new way to access Teradata Database that doesn't require a Teradata-provided client library or driver.

Teradata Customer Interaction Manager

Teradata Customer Interaction Manager (CIM) helps put into action the analytics and insights gained from terabytes of data—including campaign data, customer purchase history and trends, mobile and social footprints, or whatever other elements you want to incorporate. CIM can help organizations:

- Segment target audiences into meaningful groups.
- Visualize the impact segmentation strategies have on engagement.
- Select the right message for the right customer.
- Predict customer responses to strategic offers based on historical data.
- Send personalized, relevant and timely messages based on specific attributes and behaviors.
- Create and manage custom offers that align with a holistic brand messaging strategy.

Teradata Real-Time Interaction Manager

Teradata Real-Time Interaction Manager (RTIM) helps organizations respond to data-derived insights appropriately and in real-time. In the consumer space, RTIM makes it easy to reach customers on many different channels at any time and respond with contextually relevant messages that meet real-world needs. RTIM can help organizations:

- Make real-time marketing decisions to improve marketing strategy.
- Provide contextually relevant content for customers through any channel, online or offline.
- Tailor customer experience management to improve over time.
- Increase customer satisfaction and loyalty throughout the entire life cycle.
- Align inbound and outbound marketing initiatives with corporate strategies.
- Provide authentic individualized, one-to-one interactions with every customer.
- Automatically learn overtime the messages and channels that perform the best with each of your customers.

Teradata Integration with Microsoft Power BI

Teradata integration with Microsoft Power BI gives organizations the ability to connect to hundreds of data sources, simplify data preparation, and drive ad hoc analysis. Power BI can be used to produce reports, then publish them for an organization to consume on the web and across mobile devices. Power BI also makes it easy to create personalized dashboards with a unique, 360-degree view of their business. Power BI scales across the enterprise with governance and security built-in.

Business Connectivity

The high value of IOT-derived analytics and insights are such that organizations would do well to make data-driven insights available securely across business systems running on Teradata Database, using the Teradata REST API. Organizations can further push the value of insights through enabling secure access from mobile devices.

Real-World Examples

From an IT perspective, it can be reassuring to see that Teradata technology for pulling value from IoT and other forms of big data is well-proven and already helping organizations in the real world. Here are two examples of such firms:

An International Manufacturer of Trains

An international company that, among other industrial pursuits, manufactures railway engines, cars, and systems, uses IoT information, processed using the Teradata Unified Data Architecture (UDA) including a Teradata Database, Teradata Aster Analytics, and an appliance for Hadoop.

The company applied the Teradata Aster nPath® function to categorize different sensors attributing normal, high, and low values, then tracking changes. These changes revealed which sensor pattern was likely to result in engine failure. For example, the analysis showed that when the engine temperature dropped from mid to low then rose to mid value again, an engine was likely to fail in three days. These and similar finds were used to create predictive modeling that greatly reduced unplanned downtime for the trains.

The company credits Teradata Aster Analytics' handling of IoT from its sensors with allowing it to achieve benefits including:

- Increased up-time through significant reduction of un-planned downtime.
- Extension/flexibility of maintenance intervals because of better risk visibility.
- Reduced labor costs because quicker root-cause analysis, improved first-time-fix rates and other metrics.
- Improved asset utilization as railways were able to get more mileage from fewer cars.
- Enhanced planning, with streamlined supply chain management.
- Ability to provide tighter up-time guarantees and performance-based contracting.
- Increased service contract capture rate for providing higher recurring service revenue.
- Ability to offer service as a key differentiator.

An International Media Content Provider

One of the world's leading media content providers, a company that streams movies and other rich content into millions of homes around the world, has earned a strong reputation for its ability to suggest titles a viewer might enjoy based upon previous viewing history.

While this sounds straightforward enough, there is great nuance involved and subscribers are sensitive to how well the company can personalize its recommendations, a process that requires continuous learning and refining. The task demands the ability to quickly gather and reliably analyze millions or even billions of events every day found in a variety of data sources, formats and locations—otherwise known as big data. The company needed a big data analytics solution that could:

- Rapidly and reliably handle staggering workloads; it must support insightful analysis of billions of transactional events each day—every search, browse, stop and start—in whatever data format that records the events.
- Work with a variety of analytics approaches, including neural networks, Python, Pig, as well as varied Business Intelligence tools, like MicroStrategy, Microsoft Power BI, and Tableau.
- Easily scale and contract as necessary with exceptional elasticity.
- Provide a safe and redundant repository for all of the company's data.
- Fit within the company's cost structure and desired profit margins.

The company the solution it required using the Teradata Unified Data Architecture, including creation of a Teradata-powered data warehouse in the cloud. The Teradata UDA includes:

- An integrated data warehouse, which provides access a comprehensive and shared data environment to quickly and reliably operationalize insights throughout an organization.
- A powerful discovery platform provides discovery analytics that rapidly unlock insights from all available data using a variety of techniques accessible to mainstream business analysts.
- Hadoop integration to provide the means to economically gather, store and refine all a company's data and facilitate the type of discovery never before believed possible.

About the Author

Artur Borycki's career in IT spans more than 20 years in several organizations and roles. In his current role as part of Technology and Innovation Office, Artur is responsible for exploring, incubating, and enabling acceleration in the areas of big data, industrial IoT, and artificial intelligence. Artur works with customers, academia, and partners to prototype and drive innovative solutions in areas like precision medicine, smart nation, smart city, and industrial IoT.

About Teradata

Teradata empowers companies to achieve high-impact business outcomes. With a portfolio of business analytics solutions, architecture consulting, and industry-leading big data and analytics technology, Teradata unleashes the potential of great companies. Visit **Teradata.com**.

10000 Innovation Drive, Dayton, OH 45342 **Teradata.com**

11.17 EB7303