

T-Mobile challenges churn with data

T-Mobile's architecture helps it put data to use across the business.

by [Brett Sheppard](#) |

Sections

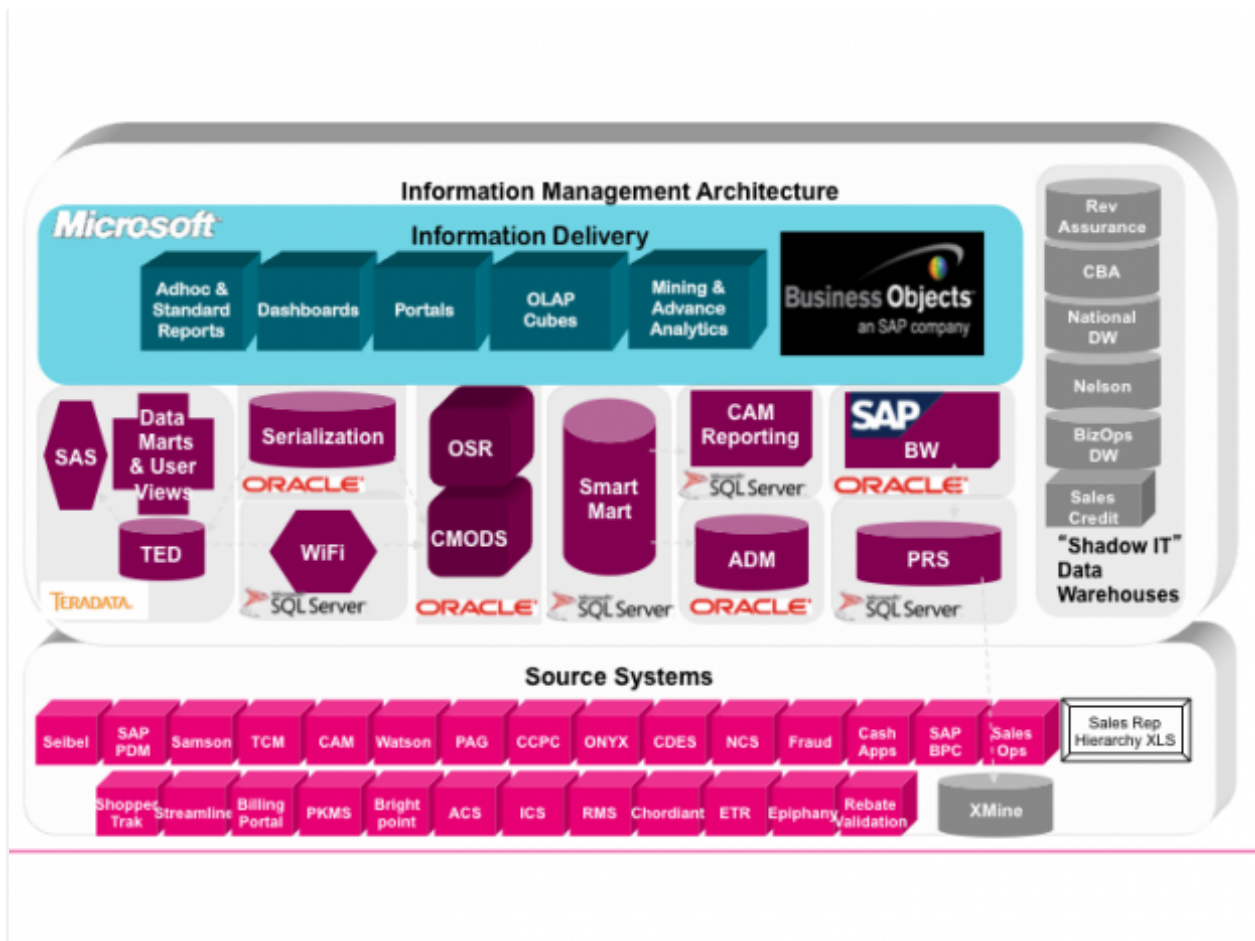
- [Information management architecture](#)
- [Virtual data zones](#)
- [Regional marketing campaigns](#)
- [Churn management](#)
- [Single-screen Quick View for customer care](#)
- [Next steps](#)

For [T-Mobile USA, Inc.](#), big data is federated and multi-dimensional. The company has overcome challenges from a disparate IT infrastructure to enable regional marketing campaigns, more advanced churn management, and an integrated single-screen “Quick View” for customer care. Using its data integration architecture, T-Mobile USA can begin to manage “data zones” that are virtualized from the physical storage and network infrastructure.

With 33.2 million customers and US\$20.6 billion in revenues T-Mobile USA manages a complex data architecture that has been cobbled together through the combination of VoiceStream Wireless (created in 1994), Omnipoint Communications (acquired in 2000) and Powertel (merged with VoiceStream Wireless in 2001 by new parent company Deutsche Telekom AG).

Information management architecture

T-Mobile USA integration architecture manager Sean Hickey and his colleagues manage data flows across a federated, disparate infrastructure. To enable T-Mobile’s more than 33 million U.S. customers to “stick together,” as the company says in its marketing tagline, a lot of subscriber and network data has to come together among multiple databases and source systems.



T-Mobile Information Management Architecture and Source Systems

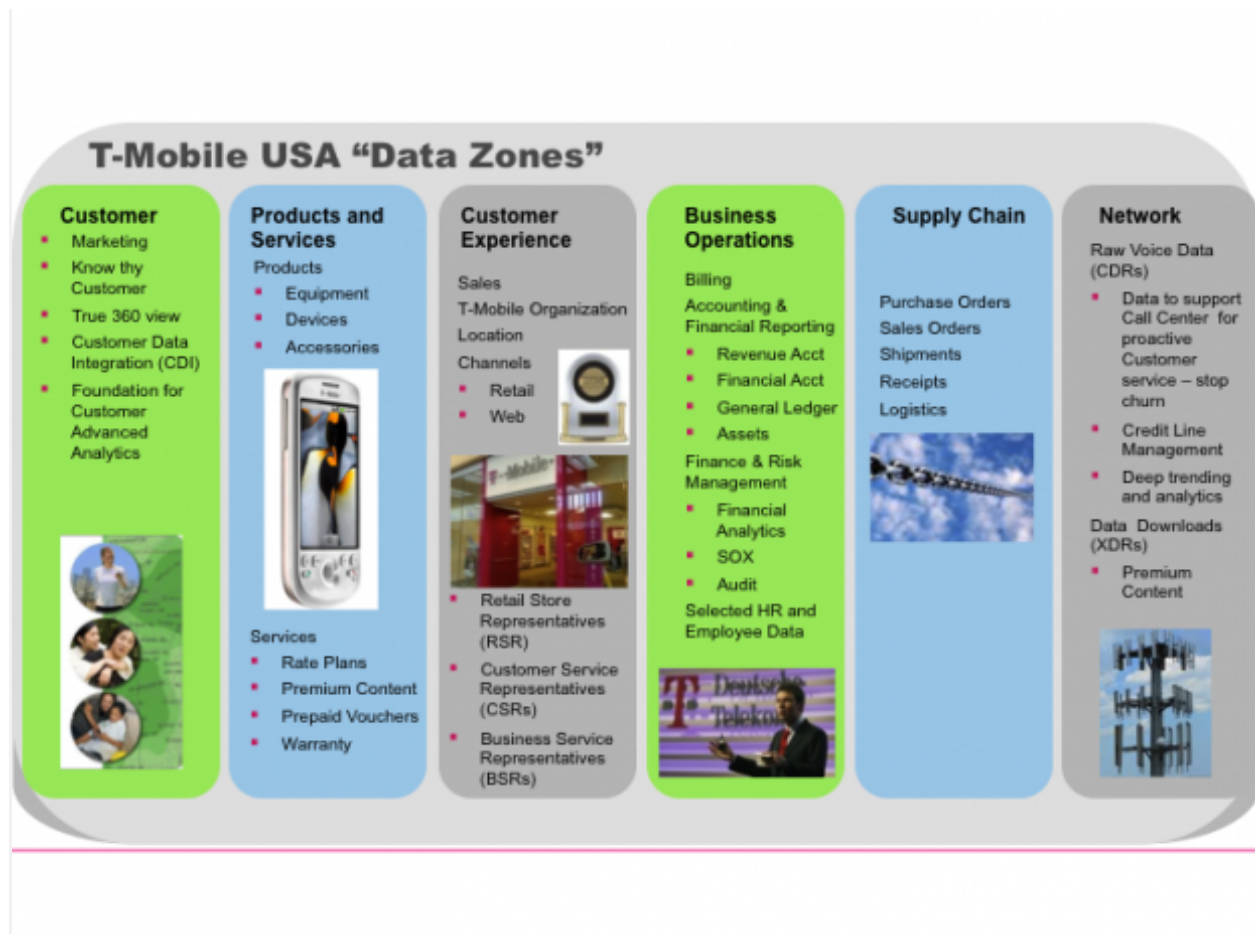
Previously, many IT systems were very point specific, stove-piped and not scalable. Some systems began as start-up projects that are now still running seven or eight years later, long after they no longer meet a good return on investment (ROI) standard. Staff that knew the original data models and schema no longer work there.

To integrate data across its disparate federated architecture, T-Mobile USA uses [Informatica PowerCenter](#). (*Disclosure: Informatica is a client of my company, Zettaforce.*) T-Mobile runs PowerCenter version 8.6.1, is a 9.1 beta customer, and plans to upgrade to version 9.1 in the fourth quarter of this year. Data modeling tools include [CA ERwin](#) and [Embarcadero ER/Studio](#). To identify data relationships in its complex IT environment, T-Mobile USA uses [Informatica PowerCenter Data Profiling](#) and IBM Exeros Discovery Data Architecture (now part of [IBM InfoSphere Discovery](#)). This data integration layer powers multiple key business drivers, including regional marketing campaigns, churn management and customer care.

To visualize data, financial analysts at T-Mobile USA use [Tableau Software](#) among other tools. T-Mobile USA uses [Splunk](#) to search across its log sources, aggregate IT data and validate PCI compliance, among other uses.

Virtual data zones

Backed by this data integration layer, the T-Mobile USA architecture team introduced the concept of virtual “data zones”. Each data zone comprises data subjects, and is tied to one or more business objectives. These zones virtualize data applications from the physical data storage and network. From a data architecture perspective, the data zone approach helps pinpoint where there are complex systems to maintain, shadow IT, redundant feeds, differences in data definitions or incompatible data. This approach also helps highlight where business rules are embedded all over, leading to duplicate or inconsistent business rules, versus more centralized rule management.

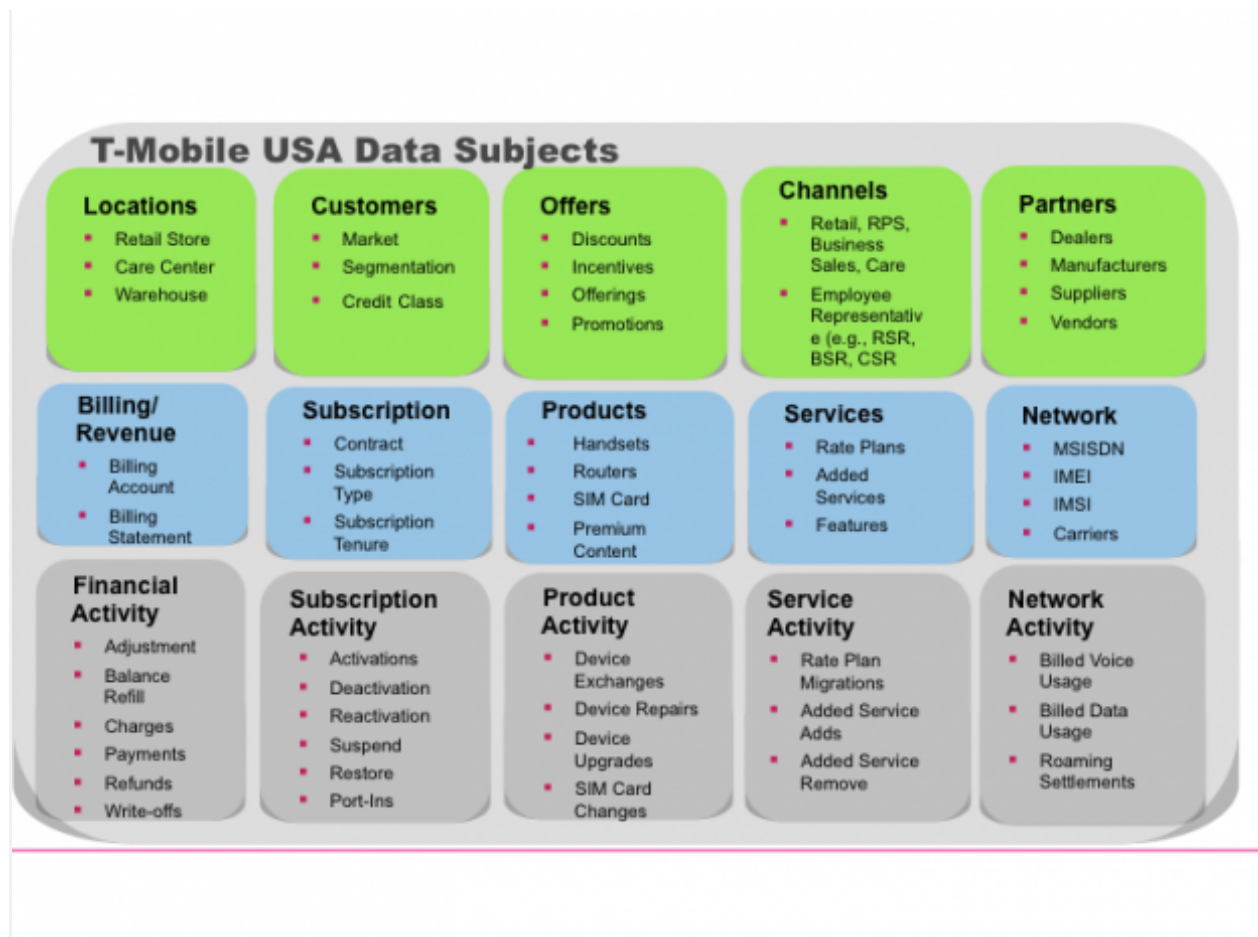


T-Mobile Data Zones

T-Mobile USA adopted [SAP BusinessObjects Strategic Workforce Planning](#), the first SAP application to use [SAP HANA in-memory computing](#) to provide real-time insights and simulation capabilities. According to Sean Hickey, T-Mobile USA has been very pleased so far with pilot tests of the HANA-enabled in-database analytics.

Legacy systems do present constraints with management of specific data subjects. For example, T-Mobile USA would like to archive off historical subscriber records that are more than seven years old, which is the cut-off date for regulatory-required retention. However, with the bottom-up growth of the the company’s data architecture, it is difficult to carve out old data. The call date was not necessarily part of the partition key. Accordingly, with how

data is segmented, T-Mobile USA continues to store subscriber records and other information dating back to 1999.



T-Mobile Data Subjects

Regional marketing campaigns

Each data zone is associated with one or more strategic business objectives. For marketing, a couple years ago T-Mobile did a fairly aggressive U.S. reorganization to become a more regional-oriented organization. T-Mobile used to do U.S. national marketing campaigns but has moved to a decentralized model that involves geography, demographics and call usage patterns to perform cross-sell and upsell campaigns by region, with assistance from third-party marketing partners for outsourced analytics. T-Mobile now has more than 20 regional districts across the United States, with a local head who is responsible for sales, marketing and operations in that district.

Northern California VP and GM Rich Garwood added about 30 staff in new regional jobs to take over functions previously handled by T-Mobile USA headquarters in Bellevue, Wash., and will for the first time make a concerted effort to market to small business owners in

Northern California. “It’s exciting for us as employees. We really have local ownership of what the results are”, Garwood told the [San Francisco Business Times](#).



T-Mobile Business Objectives Associated with Each Data Zone

[SAS Marketing Automation](#) gathers 300 attributes, including campaigns, take rates and dispositions. Before, T-Mobile did national campaigns, with a kind of “shoot and see what sticks” approach. Now, T-Mobile’s regions can run targeted campaigns specific to customer demographics and customer segmentation. This requires pulling in more than 20 different sources of data. Deep data mining operations cover billions of rows a day. For analytics reporting, T-Mobile USA uses [SAP Business Objects](#) including Crystal Reports. Finance and accounting department staff still tend to download data into Excel spreadsheets. As part of the company’s data security enforcement, every employee and contractor is required to use a T-Mobile-supplied computer with hard drive encryption. Power users can access the Teradata system directly with Teradata SQL for data mining.

Churn management

T-Mobile USA has begun using a “tribe” calling circle model — with multi-graphs akin to social network analysis — to predict propensity of churns and mitigate the potential impact of “tribe leaders” who have high influence in large, well-connected groups of fellow

subscribers. An influential tribe “leader” who switches to a competitor’s service can kick off “contagious churn,” where that leader’s friends, family or co-workers also switch.

In the past, wireless service providers calculated net present value (NPV) by estimating a subscriber’s lifetime spend on services and products. Now, part of the NPV calculation measures the level of influence and size of a subscriber’s tribe.

As [noted](#) by Ken King, marketing manager for the communications industry at SAS: “In North America, we increasingly work with service providers that are keen to examine not only segmentation, churn and customer lifetime value but new things like social networking impact on their brands, or the relationships between customers so they can recognize group leaders and their influence on others in terms of buying products or switching to competitors.”

Churn management at T-Mobile USA begins with an [Amdocs](#) subscriber billing system and financial data stored in a [Teradata](#) enterprise data warehouse (EDW). “The heart of the company is the billing system”, said Sean Hickey.

However, some key data for churn management is not captured in the billing system. Non-billable events can be very important for marketing. Raw call data gathered from cell towers and switches, supplied by [Ericsson](#) and other system vendors, can show the number of dropped calls for each subscriber and the percent of a subscriber’s total calls that drop. T-Mobile USA loads call data into [IBM Netezza](#) systems from a series of flat files. T-Mobile engineering uses this data for drop call analysis. They can look at drop calls for specific phone numbers. For example, if a T-Mobile customer moves to a new home in a location where cell towers provide only limited coverage, T-Mobile marketing can proactively offer the subscriber a new cell phone that could improve reception, or a free femtocell that connects to the subscriber’s home broadband network. Customer demographic data, however, is not stored in the Netezza systems — that’s stored by T-Mobile IT in its Teradata enterprise data warehouse (TED).

The Teradata EDW sends out extracts to T-Mobile USA’s SQL servers and Oracle servers. The Teradata, Oracle and Microsoft SQL Server databases are fed by dozens of source systems, including Seibel, the Billing Portal, Epiphany, Sales Operations and Cash Applications. “Shadow IT” data warehouses include revenue assurance, cost benefits analysis (CBA), business operations and sales credits.

The T-Mobile USA information management team has targeted multiple data marts and shadow IT warehouses to incorporate into the Teradata enterprise data warehouse, pending funded projects to add those to the EDW. In this respect, T-Mobile USA is similar to many other Fortune 500 organizations, which balance an EDW vision with the constraints of budgeting, legacy systems and acquisition integration, and therefore manage a hybrid information management architecture combining an EDW and data federation.

Data delivery is really across the board. It takes a day for information to be batch loaded from retail stores and web sales. It used to then take a second data for analysis. The combination of [Informatica PowerCenter](#) and [SAP Business Objects Explorer](#) enables the T-Mobile USA channel management team to run reports within seconds rather than an hour or a day. "It's a pretty cool platform," said Hickey. Future steps may target speeding up the data acquisition.

T-Mobile USA continues to innovate for churn management. To better identify the multi-faceted reasons behind customer turnover, T-Mobile USA ran a proof of concept (PoC) with [EMC Greenplum](#), with a storage capacity of roughly 1 petabyte, including data from cell towers, call records, clickstreams and social networks. Following the PoC, T-Mobile USA decided to work with an outsourced service provider, which uses [Apache Hadoop](#) to store and process multi-dimensional data. Sentiment analysis predicts triggers and indicators of what customer actions are going to be, which helps T-Mobile proactively respond. Informatica's newly announced PowerCenter version 9.1 includes connectivity for Hadoop Distributed File System (HDFS), to load or extract data, as explained by Informatica solution evangelist Julianna DeLua. Customers can use Informatica data quality and other transformation tools either pre- or post-writing the data into HDFS.

The advanced churn management capabilities are beginning to show results. For the second quarter of 2011, net customer losses were 50,000, compared to 99,000 in the first quarter of 2011. Further, with the combination of the churn management analytics and the integrated single-view for customer care discussed in the next section, T-Mobile USA is focused on upgrading its customers to higher quality products and concentrating on retention of loyal subscribers with a high lifetime customer value.

Single-screen Quick View for customer care

Backed by its data integration architecture, T-Mobile USA rolled out Quick View as part of an upgrade of its customer care system. With Quick View, agents and retail store associates can view multiple key indicators including the customer segmentation value on one screen. Before, call center agents and retail store associates had to look at multiple screens, which is problematic while talking live with a customer.

Quick View pops up with offers specific to that customer, such as a new phone or a new service plan. Subscribers with a high value may be sent automatically to care agents specially trained on handling high-value customers. T-Mobile USA plans to extend Quick View to third-party retailer partners such as [Best Buy](#) that sell T-Mobile phones and services in their retail stores.

Next steps

In addition to empowering innovations in regional marketing campaigns, churn management and customer care, data integration would have taken on even more significance if the

AT&T acquisition of T-Mobile USA was approved. An approved acquisition would have kicked off a host of new integration initiatives between the two companies. Instead, with the acquisition cancelled, T-Mobile USA can restart its longer term projects including ways to shift more data sets from separate data marts into the Teradata EDW along with other upcoming projects.

Related:

- [Get started with Hadoop: From evaluation to your first production cluster](#)
- [Outliers and coexistence are the new normal for big data](#)

tags: [Big Data](#), [data business](#), [data product](#), [data service](#), [mobile data](#), [stratablog](#)