INFOBRIGHT | JASPERSOFT

Austin Energy – Implementing Effective BI and Data Warehousing in Weeks

A Case Study in Using New Approaches to Solve Old Problems.

Austin Energy

Company:
Austin Energy

Ownership:
Public

Location:
Austin, Texas

Services Provided:
Electric Service Delivery, Power Production, 311 Call Center City Operations

Service Connections and Meters:
400,000+

Area Served:
437.06 Square Miles

Total Employees:
1721.75

Background

A municipal utility owned and operated by the City of Austin, Austin Energy has provided the city with electric service for over 100 years. Austin Energy is the ninth largest public power utility in the nation. Austin Energy provides power generation, distribution, transmission and substation transmission. The utility continues to provide award-winning service to its industrial, commercial, and residential customers. The Electric Service Delivery business unit of Austin Energy was the first in the nation within its industry to receive an ISO 9001:2000 Quality certification. Austin Energy's utility revenue bonds rate AA- stable from Fitch, Inc, A1 Positive from Moody's Investors Service, and AA Stable from Standard and Poors. The utility holds over $3.5 billion in assets, and generates over $1 billion in annual revenue.

The Challenge and Drivers for Change

As with most organizations for which technology is an organic evolution, Austin Energy continues to experience opportunities to provide stable, time-bound, relevant and organizationally disparate performance metrics throughout the utility. Traditional business intelligence and data warehousing models often cost millions of dollars and take years to migrate through design, development, and production-level deployment. Burdened salaries for highly-skilled IT professionals tasked with building and maintaining these systems is often at the peak of professional-services costs, and retention of these employees can prove difficult. Proprietary technologies add to costs, issues with interoperability, ongoing licensing requirements, and specialized consulting or contracting to enable a successful system. Although a request for an executive dashboard had been officially recognized as a project within Austin Energy for years, a deliverable that would meet the requestor's requirements had yet to be produced, primarily due to following the traditional business intelligence and data warehouse model.

As with most dashboarding efforts, the project goals start with a series of defined key performance indicators (KPIs) sometimes referred to as performance measures. While the underlying data that creates the utility's performance measures remains constant throughout the organization, the level of detail required for analysis varies greatly depending upon the job duties of an individual manager. For example, data that is relevant to an executive charged with health, growth and the strategies of a utility is rarely relevant to functional or process managers and supervisors, which in turn differs from the needs of individual contributors and subject-matter experts tasked with mission-critical and daily tactical operations. Traditional models for disseminating data throughout these levels of an organization require the use of data cubes or aggregate data tables. For instance, few relational database management production systems can handle multiple ad-hoc requests for summaries or averages grouped by varied and disparate filters from the raw source data.

Thus, in the traditional data warehousing model, aggregations of data are created and made available that attempt to meet the needs of these ad-hoc queries. These aggregate tables require creation, maintenance, and deletion on a constant, ongoing basis. Add to this the complexity of communicating the business requirements needed to the creators of the aggregate tables from the subject matter experts,[i] and it becomes obvious why 60-90% of traditional data warehousing efforts fail.[ii]

Austin Energy's solution required a timely, low-cost, highly available, service-oriented architecture that would enable any authorized user to analyze relevant and timely data. The key project deliverables and objectives were as follows:

1. Agile Deployment
2. Low-Cost
3. High Availability
4. Able to use existing infrastructure
5. Ability to provide simultaneous dashboards, reports and unit-level data analysis
6. Ad-hoc query capabilities for business unit subject matter experts
7. Ease of institutionalization
8. Secure, stable platform
9. Infinitely configurable
10. No requirement for specialized workforce
11. Scalable

The Solution

After years of approaching the executive dashboard project in the traditional manner, Austin Energy's project managers realized that to meet all of these seemingly polar objectives would require a non-traditional approach. In June of 2009 an evaluation of technologies that could meet the rigorous demands of the solution proved extremely successful. All of the objectives could be met through the acquisition and deployment of just two technologies, a Business Intelligence suite of tools provided by Jaspersoft, and a revolutionary column-based database engine available from Infobright. Through the use of these solutions, two project managers, working part-time over six weeks, were able to meet all of the requirements of the project and subsequently begin a staged institutionalization of the performance management solution.

Project Implementation and Deployment Strategy

The request for a configurable executive level dashboard was the first need to be addressed. Due to ISO 9001:2000 requirements for performance measures within the Electric Service Delivery business unit, production of an operations and maintenance (O&M) cost per customer key performance indicator (KPI) was paramount. Using JasperETL, Jaspersoft's Extraction, Transformation and Load (ETL) tool[iii], Austin Energy's team created automated jobs that pull raw data from three distinct sources: a nightly flat text file of all of Austin Energy's accounting journal entries, a nightly extract of Oracle data containing organization structure, and a nightly extract from an AS400 IBM CIS system that provided numbers of customers billed by period. For this KPI, two years of journal entries constituted 2.5 million rows of data in 73 columns.

Once these automated jobs were validated and running successfully, the output of a relational data model into the Infobright® engine running on MySQL allowed for extremely impressive data compression and response time. In a traditional database, two and a half million journal entries with primary and foreign keys, partitioned and indexed, require nearly five gigabytes of storage; additionally, creation of aggregate tables to provide acceptable response times in dashboards and reports is required. Using Infobright technology, storage requirements went from nearly five gigabytes to 74 megabytes, (a compression factor of greater than 98%). More importantly, the technology allows for analysis at the line-item level, negating the need for aggregate tables. By not having to create aggregate tables to provide dashboards and reports, development time to produce these deliverables was reduced by a factor of five.[iv] Equally important, because jobs were not required to build aggregate tables and traditional data cubes, maintenance of the system was dramatically reduced while reliability was improved.

The ease of use of this system, combined with its seamless interoperability with existing and proprietary architectures, allowed job creation, testing and deployment within weeks.

The project now had a data warehouse that updated and validated automatically. Availability of financial data reduced data age[v] from up to six weeks to over night, a factor of 1/42nd. Using this new data warehouse, creation of dashboards and reports was simply a matter of drag-and-drop with iReport, Jaspersoft's report designer. These reports are incorporated into a secure dashboard, using Jaspersoft's JasperServer that also allows the user to do ad-hoc queries. In addition, access and authority rights provided by Active Directory govern the system[vi]. This last item is important to note. By simply tying in to the existing organizational structure, the project was able to provide relevant information at the user level: Executive Dashboards were available to only those users with rights to view them, while reports and ad-hoc data analysis became available to individual contributors and supervisors/managers with their own corresponding data access privileges.

A primary value of a data repository is the reduction in time required for report generation. A monthly O&M Detail report for the entire utility that previously took ten person-hours to generate could now be created instantaneously with a mouse-click. Moreover, the subject matter expert charged with creating the report now had the ability to generate similar reports, *without having to contact a specialized foreign department or group*. By creating a data model and repository accessible by the entire organization as per the security requirements, the project put the data in the hands of the business experts.

One added benefit of the initial phase of this project was the availability of data. Because the accounting data encompassed the entire utility, the executive dashboard, detail reports and query analysis created originally for one business unit, Electric Service Delivery, now provided O&M cost data to *all business units simultaneously.* Had the project continued with a traditional data warehousing model, a year to produce one metric for one business unit could easily be extrapolated into one metric per business unit per year. With eleven business units each with ten KPIs, it's hard to imagine that the utility would wait for over a century for relevant metrics. In contrast, Austin Energy now had one metric available to eleven business units in weeks.

As if the success of this project doesn't already seem too good to be true, there's one last item of note: costs for development time, software procurement and ongoing maintenance is a factor greater than one-tenth that of traditional business intelligence and data warehouse models. When combined with the reduced cycle time and additional benefits meeting the aforementioned project deliverables, the return on investment is significant.

Lessons Learned and Future Plans

As Austin Energy matures using the new data warehousing and reporting model, the remainder of the financial data, as well as more specific electric service delivery data, will be made available. Human resources, customer care, and other non-utility specific business areas are benefiting from the solution, and have requested similar projects to meet their reporting, scorecard and data analysis requirements. Additionally, as the ease of the business intelligence tools demonstrates fluid socialization, end users within the business channels will generate more reports.

No software or hardware solution will replace the need for superlative requirements analysis and project management. However, the greater the complexity of systems, applications and requirements, the greater the potential for delay, scope reduction, or outright failure of an attempted solution. Using two products that work together to provide an end-to-end suite for business intelligence and data warehousing greatly reduces complexity.

From a socialization perspective, it will benefit future projects to have robust hardware architecture in place that is stable as migration from test instances to production instances occur. As with any technology an increase in concurrent users has a direct negative impact to response time and functionality. From the end-user's standpoint, it is imperative that response time or functionality show little to no latency. Otherwise, the tendency to blame the software or solution design becomes prevalent.

As a public utility, Austin Energy has an opportunity and obligation to reduce costs of service. This in turn allows for maintaining low rates to customers. Any advantage gained through automation, technological improvements, and low-cost solutions should be awarded preferential adoption. Although the application of this solution demonstrated significant savings in time, cost, effort, and maintenance requirements, it is not industry-specific. The products, principle and performance of this project have wide-ranging benefits for any organization or institution.

[i] What data architect intuitively knows the difference between encumbrance, SAIDI, GIS coordinates, or customer satisfaction ratings?

[ii] Voelker, M. (2001). Databases: The next generation. Insurance & Technology, 26, 3, 30-34.

[iii] Licensed from Talend

[iv] This factor does not take into account the avoided cost and resource burden of maintaining additional tables and structures.

[v] the time it takes for data to become available once created

[vi] Jaspersoft has interoperability with many LDAP solutions. Austin Energy uses Active Directory.

JASPERSOFT

Jaspersoft Headquarters
539 Bryant Street, Suite 100
San Francisco, CA 94107
1-888-399-2199
Phone: 415.348.2380
Fax: 415.281.1987
Email: sales@jaspersoft.com
jaspersoft.com

INFOBRIGHT

Infobright Headquarters
47 Colborne Street, Suite 403
Toronto, Ontario M5E1P8
Canada
Tel. +1 416 596 2483
Toll Free +1 877 596 2483
Fax +1 416 352 5674
infobright.com