

Using MySQL 5.1 for Data Warehousing / Business Intelligence

Overview

Breadboard BI provides affordable business intelligence (BI), data integration, and general software application development consulting. We also offer a suite of customizable BI modules that form a modular and scalable enterprise solution. Our mission is to provide products and service that enable companies to reduce costs, improve sales, and satisfy regulatory requirements.

We have experience with a full arsenal of lower-cost open source tools like MySQL and Pentaho, as well as proprietary technologies. Our talented staff includes business intelligence architects, enterprise programmers, database administrators, and systems administrators. Our combined expertise includes most functional business areas, open source and proprietary tools, application servers, databases, and operating systems. Whether it's affordable business intelligence systems or web application development, our staff is ready to build it. A collection of Business Intelligence examples can be evaluated on our demo server. The backend database for all of these examples is MySQL Enterprise 5, the front-end is using Pentaho.

At Breadboard BI, we have used MySQL 5.1 as the backend to several data warehouse/business intelligence solutions. We have recently completed a very exciting project on behalf of Boats.com / YachtWorld.com. This project is based on a 100% open source stack (Linux/MySQL/Pentaho) and performs very well. On the MySQL side, we have utilized the MyISAM engine, table partitioning, and an aggregate table strategy to support a presentation layer of reports and dashboards. We would like to thank MySQL for creating such a great database, it allows small companies like Breadboard to deliver great solutions for very little money.

Business

Together, YachtWorld.com and Boats.com are the single most powerful and effective global marketing solution for the boating industry. With over 140,000 new and used boat listings worldwide offered by over 4,000 brokers, dealers and builders in 115 countries, YachtWorld.com and Boats.com are visited by over 4 million boating consumers every month who click through over 95 million page views.

YachtWorld.com is the premier online sales channel for yacht brokers around the world, providing a complete suite of online marketing services through BoatWizard(tm), its proprietary back-end tool. Boats.com provides marketing and Web services to new boat dealers and builders, and offers a "For Sale By Owner" classified service. Headquartered in Seattle, Washington, YachtWorld.com and Boats.com have European headquarters in the United Kingdom, with sales offices in

Germany, Italy and Russia, and sales representation in Dubai, Australia and China. Both companies are business units of Dominion Enterprises, based in Norfolk, VA.

Challenge

Like many companies, YachtWorld.com and Boats.com were collecting large quantities of data from their diverse operational systems – millions of daily page views from each of their popular web sites, as well as significant volumes of sales lead emails and toll free calls to their affiliated brokers, dealers, and builders.

Understanding the potential value of this data, a project was undertaken to build a market intelligence product. However, unlocking the valuable information in this data was a real challenge due to the large and disparate data sources. The web-related data resided in very large log files, the call and email data was distributed across multiple, heterogeneous MySQL and Oracle databases, other supporting data was found in various flat files. YachtWorld.com and Boats.com required a flexible, highly-scalable, maintainable market intelligence solution that could bring together all the data, and present the information in an easy to use format. The final challenge - the solution needed to be built within a reasonable budget.

Solution

YachtWorld.com and Boats.com required a flexible, highly-scalable, maintainable market intelligence solution that could bring together all the data, and present the information in an easy to use format. The envisioned market intelligence product required a centralized data warehouse from which users could perform sophisticated analysis based on cubes, dashboards, and reports. Both open source and proprietary database and business intelligence solutions were considered during the selection process. MySQL was chosen as the database engine; the Pentaho business intelligence suite was chosen to provide extract, transform, load (ETL) and analysis/ reporting functionality.

Both selections were based on their advanced functionality (especially MySQL 5.1 partitioning); large, active user-bases; existing internal expertise with the engine (MySQL); the long-term viability of the companies; and the availability of professional support subscriptions. Breadboard BI, Inc. provided the project management and consulting services necessary to architect and build the solution. They were chosen based on a proven track-record of delivering scalable business intelligence solutions, as well as their MySQL and Pentaho partnerships and certifications.

The new YachtWorld.com and Boats.com solution, named BoatsReports, represents a 100% open source stack - from the Linux operating system on the servers, to the MySQL database, to the Pentaho data integration and presentation layers. This stack required zero upfront investment in database or business intelligence software

licenses (usually running well into USD six figures), making it easier to receive internal project approval. In addition, the open architecture of this stack made it easy to integrate with existing systems. Although a contingency plan was developed to swap out open source tools for proprietary or custom- developed solutions, these alternate plans were not needed.

Results

Boats.com and YachtWorld.com have affordably transformed their oceans of disparate data into a comprehensive market intelligence product based on open source technologies. The actionable information in BoatsReports offers internal analysts and external partners previously unavailable insights. Specific examples of advantages now offered to boat manufacturers include the ability to:

- Accurately monitor their sales leads and sales leads of their participating dealers by model and location;
- Know which boat models are generating the most interest with the buying public in different countries; and
- Understand how their boat models measure up against the competition within each market segment.

Making use of previously unavailable information with the help of open source tools, this powerful solution has quickly proven invaluable in the decision making process for production, sales and marketing strategies. Because of the size of the Boats.com & YachtWorld combined data set, the data in aggregate form now represents some of the most comprehensive benchmark data in the boating industry - by itself a valuable commodity allowing for unmatched competitive analysis. A secondary result of the new solution has been that it has exposed data issues in the source systems. Several projects have been undertaken to further improve operational data, allowing customers to be better-served.

SOURCE-MySQL website

DISCLAIMER: The content provided in this article is not warranted or guaranteed by DBA Guys. The content provided is intended for entertainment and/or educational purposes in order to introduce to the reader key ideas, concepts, and/or product reviews. As such it is incumbent upon the reader to employ real-world tactics for security and implementation of best practices. We are not liable for any negative consequences that may result from implementing any information covered in our articles or tutorials. If this is a hardware review, it is not recommended to open and/or modify your hardware.