

MCLAREN ELECTRONIC SYSTEMS

Motorsport innovator investigates Parallel Data Warehouse (PDW) capabilities for analyzing huge volumes of time-series data



ESSENTIALS

Challenge

- To extend and enhance the capabilities McLaren Electronic Systems offers its customers in leveraging realtime acquisition of data for improved innovation and optimization

Solution

EMC Consulting:

- Built a data warehouse solution based on the Microsoft SQL Server 2008 R2 Parallel Data Warehouse (PDW) appliance
- Developed analytics that could leverage MPP performance of read-intensive workloads against very large data sets to perform analyses on hundreds of billions of rows of historical data

Key benefits

- All data from a single session is available for analysis in less than seven minutes
- Analytics can be performed across an entire set of historical data—rather than limited subsets—facilitating the ability to identify, search for, and match particular patterns of fluctuations to learn more about their cause and effect as well as predict and prevent problems

McLaren Electronic Systems develops and manufactures advanced electronic control systems and software for customers in motorsport, automotive, aerospace, and healthcare industries. Together with technical partner Microsoft®, the company is the Official Engine Control Unit (ECU) Supplier to the FIA Formula One World Championship.

In a sport where the slightest edge of innovation or optimization can mean the difference between winning and losing, one of McLaren's leading products, SQL-Race software, helps facilitate that advantage by managing the realtime acquisition of data from racing cars. A highly compressed storage engine built on a Microsoft SQL® Server database, SQL-Race makes possible a realtime, easy-to-interpret graphical representation of data on McLaren's Microsoft Windows®-based ATLAS display and analysis system.

With the release of Microsoft's SQL Server 2008 R2 Parallel Data Warehouse (PDW) platform, McLaren recognized an opportunity to extend and enhance the capabilities it offers its customers, and raced to be one of the first companies in the world to build a massively parallel processing (MPP) Microsoft SQL Server data analysis application to analyze huge data sets.

PDW is an appliance-based solution that uses software from Microsoft and hardware components from multiple vendors, including EMC. It provides a highly scalable, massively parallel processing (MPP) architecture that delivers high performance and incorporates Microsoft SQL Server business intelligence features and functionality.

Because MPP architectures are modular, they are able to efficiently scale to hundreds of terabytes of data and perform analyses across trillions of rows of data. PDW applications can scale to deliver the query performance required, because applications run in parallel on independent servers with separate operating system instances. In addition, because extra processing and memory power and operating system instances are added along with storage, MPP solutions scale very linearly.

"McLaren Electronic Systems builds all of its off-car applications on Microsoft platforms, and has done so for over a decade," says Peter van Manen, managing director. "The company's software solutions already provide a highly efficient way of handling large quantities of continuous data in real time. PDW offers the chance to enhance and extend this approach even further."

THE CHALLENGE

SQL-Race leverages Microsoft SQL Server 2008 to store time series data and distribute and query high frequency realtime data streams. A Formula One car generates hundreds of millions of data points per hour, with the system monitoring and measuring the performance of thousands of car components hundreds of times per second. During one race or test run, a single car can produce upwards of 700 million data values, generated over a period of less than two hours.

Over the course of a season, race teams collect many hundreds of billions of measurements. Retrieving, summarizing, and mining such huge sets of historical data to derive meaningful and useful results is an exceptionally difficult data management challenge for motorsport teams. It also poses a significant challenge for McLaren clients in manufacturing, healthcare, and other industries, who need to analyze high volumes of historical time-series data from monitors, sensor networks, and streaming data sources.

PRODUCT KNOWLEDGE AND INTEGRATION EXPERTISE YIELD A SUCCESSFUL SOLUTION

McLaren engaged EMC® Consulting to build a data warehouse solution on a Microsoft SQL Server 2008 R2 Parallel Data Warehouse (PDW) appliance. EMC Consulting work included developing analytics that could leverage MPP performance of read-intensive workloads against very large data sets to perform analyses on hundreds of billions of rows of historical data.

In addition to prior positive experiences, McLaren chose EMC Consulting again because of its Microsoft solution partner status and the fact that it offered EMC CLARiiON® AX-4 series storage for the Dell PDW appliances.

“We had worked side-by-side with EMC Consulting to develop SQL-Race two years ago,” says van Manen. “We knew they understood our application and Microsoft technologies. Because of their close partnership with Microsoft, they were also willing to work with Microsoft beta technologies such as FILESTREAM, and they were able to build us an early PDW appliance. We also wanted to leverage their experience and expertise on large data warehouse and business intelligence projects.”


In just a little over three weeks, EMC delivered a complete end-to-end PDW solution with business intelligence tools leveraging SQL Server 2008 R2 software. The EMC consultants developed extract-transform-load (ETL) scripts to load data onto the new system, created and populated tables, designed the indexing, wrote queries, and created sample reports.

EMC consultants then took advantage of SQL Server 2008 R2 Parallel Data Warehouse data distribution and parallel processing features to tune workloads and optimize query performance. They remodeled the base time-series data to support historical analysis by enabling data to be distributed across nodes and partitioned. Near perfect distribution was achieved using the Timestamp column.

The PDW appliance used for the test was based on Dell PowerEdge servers and EMC CLARiiON AX-4 series Fibre Channel-attached storage arrays, and consisted of a control rack and a single data rack with eight database nodes.



EMC Consulting leveraged a highly scalable Microsoft SQL Server 2008 R2 Parallel Data Warehouse (PDW) appliance using EMC CLARiiON AX-4 storage and Dell PowerEdge servers to analyze huge quantities of telemetry data.



THE RESULTS

EMC loaded 1.5 TB of session data into the PDW environment, achieving speeds of over 100 MB per second. This enabled all of the data from a single session to be available for analysis in less than seven minutes. Data was then scaled up to 12 TB to simulate 90 complete race and test sessions, allowing cross-session pattern analysis against almost 400 billion rows of telemetry data.

EMC Consultants were then able to use a pattern of values from one session to search for similar patterns in others, with average query times of 13 seconds, from the full 400 billion row data set.

“We wanted to be one of the first companies in the world to take advantage of analysis using Microsoft’s massively parallel processing data warehouse. In close collaboration with Microsoft, EMC Consulting helped us build a PDW appliance, write ETL scripts, load the data, and develop analytics built on our software in just three weeks. This complements our SMP-based solutions with a proven PDW reference architecture to enable new kinds of historical and ad hoc analyses of huge data sets at an affordable price.”

PETER VAN MANEN, MANAGING DIRECTOR, MCLAREN ELECTRONIC SYSTEMS

These tests demonstrated that with Parallel Data Warehouse, it was possible to do analytics across the entire set of historical data, rather than limited subsets. Having the capacity to store and analyze huge volumes of data with adequate performance is compelling. Parallel Data Warehouse allowed interactive analysis and enabled questions such as “Where has this happened before?” and “How does this relate to...?” to be answered in seconds. Such a capability could be used for identifying, searching for, and matching particular patterns of fluctuations (such as changes in pressure, temperature, or stress on particular components) to learn more about their cause and effect, as well as to predict and prevent problems.

“With data sets continuing to grow in size, and ever increasing levels of scrutiny of the interactions and influence of different parameters on performance and reliability, MPP is set to become an important tool in the armory when it comes to extracting more value from what we measure and simulate,” says van Manen.

PDW integrates fully with existing Microsoft SQL Server Business Intelligence tools. Part of the technology incorporated into PDW includes a Parallel Data Copy capability which enables “hub-and-spoke” architectures to be rapidly deployed. A centrally managed “hub” contains detailed enterprise data-feed-dependent spokes such as Fast Track data warehouse implementations, SQL Server 2008 enterprise data warehouses, and SQL Server 2008 Analysis Services OLAP databases. EMC Consulting was able to make use of the Parallel Data Copy capability to deliver a data mart to an external SQL Server 2008 R2 instance. SQL Server Reporting Services linked data from SQL Server 2008 R2 and PDW to provide session-based geospatial reports, which added a graphical element to core complex, cross-session analytics delivered by PDW.



CONTACT US

To learn more about how EMC products, services, and solutions help solve your business and IT challenges, contact your local representative or authorized reseller—or visit us at www.EMC.com.

EMC², EMC, CLARiiON, the EMC logo, and where information lives are registered trademarks or trademarks of EMC Corporation in the United States and other countries. All other trademarks used herein are the property of their respective owners. © Copyright 2010 EMC Corporation. All rights reserved. Published in the USA. 11/10 Customer Profile H7480

EMC Corporation
Hopkinton, Massachusetts 01748-9103
1-508-435-1000 In North America 1-866-464-7381
www.EMC.com

EMC²
where information lives[®]