

Smarter Analytics with IBM i for Business Intelligence

March 6, 2012

Doug Mack

Companies worldwide invest millions of dollars in operational applications to improve the way they conduct business. While these systems provide significant benefits, clients often struggle to meet the needs of various users with analysis of the data collected and stored in those systems.

Analytics is a broad term relating to applications designed for purposes of analyzing data. Key to analyzing that data is providing fast, easy access to that data while delivering it in formats or tools that best fit the needs of the end user. Business Intelligence (BI) applications are a subset of “analytics” and provide visualizations of that data for end users. Data Mining or “predictive analytics” tools provide insight into data relationships through statistical algorithms to identify patterns in the data that would otherwise go unnoticed. Data Warehouse, Data Mart or Operational Data Store architectures can be utilized to provide an optimized and isolated environment for the analytical applications.

Improved data analysis can provide many benefits including:

- Ability for end users to get the data when they need it and in the forms they need it in through a self service model.
- Executives can monitor the state of the business through Key Performance Indicator dashboards
- Spot trends and exceptions in the data with real time analytical processing
- Automate the creation and delivery of static reports in PDF, spreadsheet, or formats that allow data to be further analyzed in smart mobile devices.

Armed with timely, intelligent information that is easily understood (because it is delivered in business terms) the Business Analyst is enabled to affect change and develop strategies to drive higher profits.

At the core of any analytics solution is a database management system capable of processing complex queries in an optimal fashion and a Business Intelligence application that delivers that data to the business analysts when they need it, in the forms they need it in! More and more, business intelligence is becoming a mission critical part of every day’s activities. This paper focuses on how IBM i and the DB2 for i database deliver the key attributes required for analytics.

Why use IBM i on Power Systems for Analytics?

As the importance for understanding data becomes more and more critical for companies in today’s competitive landscape, the requirement for having a reliable infrastructure that is robust, scalable, optimized for query processing - without sacrificing security - becomes paramount.



DB2 for i is an advanced, 64-bit Relational Database Management System (RDBMS) that leverages the high performance, virtualization, and energy efficiency features of IBM’s Power Systems. A member of IBM’s leading edge family of DB2 products, DB2 for i provides a secure, reliable and scalable platform tightly integrated with the operating system, IBM i, resulting in lower total cost of ownership with its unique autonomic computing (self-managing) features.

Keep it simple. Why reinvent the wheel for your BI applications when you can leverage your computing and support infrastructure that already exists around your IBM i production environment? Leveraging an existing infrastructure allows you to “do what you do best” – reducing costs and risk by supporting your BI environment using existing administrative policies and procedures, such as security policies, backup and disaster recovery environments, and common procedures such as installation and software update processes.

The secret sauce: Integration! The cornerstone to reliability and low total cost of ownership of IBM i on Power Systems is the complete integration of database, operating system, and server. The tight integration of DB2 and IBM i provide unique benefits that allow you to concentrate on running your business while maintaining the mainframe-like reliability and security that mid-market companies demand as well. Extend that benefit into Business Intelligence by leveraging the tightly integrated DB2 Web Query for i product set that is designed from the ground up to leverage the attributes of DB2 for i that customers have realized for years.

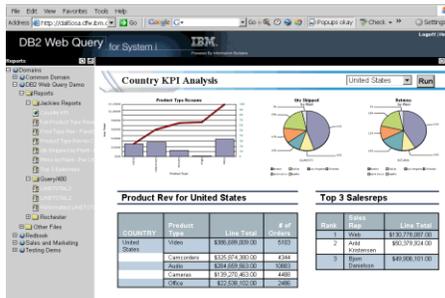
Keep the data in DB2. With DB2 for i, it's already there as the core element of the IBM i operating system. DB2 for i leverages the virtualization and large scaling capabilities of the IBM Power System platform, including Dynamic Logical Partitioning, Capacity Upgrade on Demand, and PowerVM virtualization to simply and quickly respond to changing workloads to ensure business continuity in a dynamic analytics environment. DB2's sophisticated cost-based query optimizer takes advantage of the IBM i's unique single level storage architecture, large memory addressability, and I/O and symmetric multi-processing parallelism to deliver the performance required without complex data and/or multi-server partitioning. DB2 for i also leverages the large memory configurations of Power Systems with in-memory database technologies to speed up analytics processing.

The following provides more details around these key attributes that makes the DB2 for i on Power Systems platform ideal for an analytics infrastructure!

- **DB2 Web Query for i makes analyzing DB2 for i data EASY**

IBM's DB2 Web Query for i provides capabilities to query or build reports against data stored in DB2 for i databases through browser based user interface technologies. Build new reports with ease through the intuitive Info Assist graphical report builder. Simplify the management of reports by leveraging parameterized reporting to reduce the number of reports that require maintenance. Deliver data to end users in many different formats, including spreadsheets, PDF, HTML or powerful web visualizations. Automate report execution and distribution with job scheduling and e-mail integration functions. Enable BI as part of the day to day process by integrating reports into line of business applications through the ability to embed report execution calls directly from traditional 5250 or web based applications.

DB2 Web Query is designed to enable you to modernize your existing Query/400 environment. Leverage OLAP reports to provide data slicing and dicing and trend analysis. A single OLAP report could replace many Query/400 reports required to delve into the data to answer questions from the analyst. View reports visually with the over 100 supported graphs and charts. Place several reports on a single view for an executive level dashboard. Support users on the go with the automated report distribution and mobile computing support.



DB2 Web Query accesses data in DB2 for i through a native adapter, allowing it to take full advantage of the sophisticated query optimization technology that IBM has been investing in for over a decade. Build reports seamlessly from local or remote DB2 for i databases, or even Microsoft SQLServer databases with the optional SQLServer adapter.

Key to any true BI tool is a meta data layer that shields report authors and end users from the complexities of the back end database while providing consistency of data definitions across the enterprise, eliminating ambiguity and distrust of the data in reports. The meta data layer positions DB2 Web Query as way more than just a graphical reporting tool by extending the number of personnel that are capable of building reports and eliminating the dependency on that one person who is the database expert!

DB2 Web Query Standard Edition combines all of the core features and products of DB2 Web Query into a single package supporting 100's of run-time users and a very robust set of BI capabilities.

- **SIMPLIFY the BI infrastructure with an optimized IBM i for BI Solution**

Customers are often looking to enhance the availability and isolate the BI workload from production system applications by leveraging a 2nd server. IBM i for Business Intelligence is a packaged solution that is easy to order and easy to implement, a perfect choice for improving your capability to analyze data – turning it into information that can help transform your business.

This solution combines the strengths of Power 720 Express systems, DB2 for i, DB2 Web Query for i, and InfoSphere data transportation software to deliver an integrated platform that houses the extracted and transported data.



The Power 720 Express with IBM i provides a technology foundation with proven reliability and security for the small or mid-sized company seeking a complete, integrated business intelligence system to avoid increased spending and staffing requirements while becoming more responsive to your customers, improving productivity and keeping data secure. EnergyScale™ Technology offers Intelligent Energy management features, which can dramatically and dynamically conserve power and further improve energy efficiency. These Intelligent Energy features enable the POWER7 processor to operate at a higher frequency if environmental conditions permit, for increased performance and performance per watt; or alternatively, operate at a reduced frequency if user settings permit, for significant energy savings.

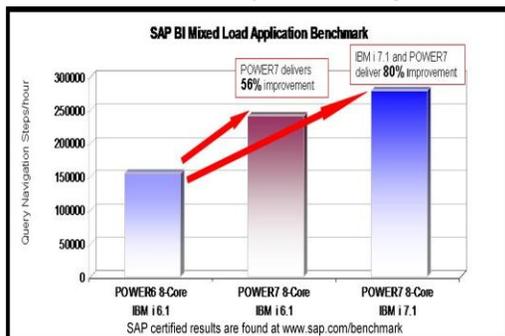
Demonstrate immediate results by leveraging the 3 days of installation services that are bundled into the price of the solution. Each of the preconfigured POWER7 based systems comes with additional processor capacity in them which can be quickly enabled for growth or additional purposes. Expand the use of the system into a full blown data warehouse or leverage it for other purposes, such as disaster recovery or creation of a development or test/QA environment.

- **DB2 for i Query Processing PERFORMS!**

IBM has been enhancing DB2 for i over the last decade, and continues to add new capabilities to enable the most efficient processing of the BI workload while still maintaining the low total cost of ownership value proposition.

The sophisticated cost based **query optimizer** built into DB2 for i is the basis of achieving optimal performance in a BI application. The optimizer provides the “brains” behind the SQL processing which most BI applications depend on. The optimizer’s goal is to build the best plan for accessing data. DB2 for i’s cost-based optimizer uses statistics collected and stored automatically in the database as well as other information, such as system configuration, parallelism settings, and available indexes in its development of a data access plan.

A great proof point of the performance capabilities of DB2 for i running a Business Intelligence workload is a recent comparison of SAP certified benchmarks run on different levels of IBM i and Power Systems. The chart below shows the most recent results of the SAP Mixed Workload BI benchmark comparing IBM’s new POWER7 system running the latest DB2 for i version (7.1) to the previous levels of Power systems and operating systems. The 80% performance improvement with IBM i 7.1 on POWER7 depict the huge performance benefit gain from the latest IBM innovations!



The 80% performance improvement with IBM i 7.1 on POWER7 depict the huge performance benefit gain from the latest IBM innovations!

Another unique advantage of DB2 for i is IBM's patented **Encoded Vector Indexing (EVI)** query accelerator technology, built into DB2 for i. EVIs use a symbol table and vector array indexing technology that can make complex BI queries scream! Many customers have reported that through EVI usage they have been able to reduce query jobs from hours to minutes, or in some cases, even to seconds. EVIs can be built right on top of your DB2 databases, and maintained real time – no need to move data to a specialized query accelerator database. With DB2 for i 7.1, EVIs can also contain aggregates, or summary data, that can also significantly reduce processing time for some queries.

Additional enablers for processing complex query workloads include:

- Sophisticated processing of star schema or snow flake data models

Star schema and snowflake data models are very common in BI implementations. The DB2 for i smarter optimizer combines a number of techniques to add special processing of queries with this type of data model. The optimizer can intuitively analyze the FACT and DIMENSION tables that make up the star-schema and improve performance by minimizing I/O (the slowest part of the performance equation)! A query can be re-written on the fly, and EVIs and dynamic bit-mapped indexing take over to process the query.

- Auto recognition of summary data with Materialized Query Tables (MQTs)

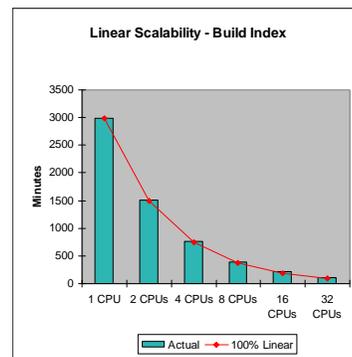
MQTs are “optimizer aware” summary tables. Building of summary tables often speeds up commonly run queries that are typical in a BI environment as the analyst often wants to assess data at a summary level. DB2’s ability to **re-write queries** allows the database engine to automatically redirect an end user query to use the MQT for faster processing.

- Autonomic Functions

Autonomic is a term associated with a self learning, self adapting organism. With DB2 for i 7.1’s Adaptive Query Processing (AQP) the query optimizer can make real-time adjustments to its execution plan, such as changing the join order or utilizing a new index, while the SQL request is running. These adjustments are implemented without any disruption to the application. Autonomic Indexes were first delivered in IBM i 5.4, and allows DB2 to create indexes behind the scenes to improve query performance while preserving those indexes for use by other jobs and users.

• LINEAR Database Scalability

A key consideration in any database platform in support of BI applications is its ability to scale while maintaining acceptable performance. DB2 for i offers a variety of techniques to address scalability/performance. **Parallel database** operations through the **DB2 Symmetric Multiprocessing (SMP) feature of IBM i** provide the ability to split single database tasks/requests across multiple processors. The SMP feature provides significant performance enhancements to database queries, index builds, and other database tasks invoked from the BI application.



The Parallel I/O design of IBM i uses intelligent I/O Processors with their own levels of caching and parallelism. Combined with Single Level Storage, the I/O subsystem of the IBM i removes a common bottleneck with data warehouse workloads: the I/O subsystem.

The graph here represents an actual test of linear scalability. This particular database task is an index build, but it shows that each time the job was run with the number of CPUs doubled, the performance indeed nearly doubled, reflecting a “nearly 100%” linear scalability.

“We built identical models over identical sets of data. There was no comparison in performance. In fact, the bigger the tables, the bigger the advantage of DB2 over SQL Server. That proved to us that DB2 for i was a better choice.” – Sam Gottlieb, Project Lead, Elie Tahari

• Low Total Cost of Ownership

IBM i’s unique architecture, including Single Level Storage and a tightly integrated object based operating system, provide the foundation for its leadership in autonomic computing. DB2 for i leverages those architectural attributes to reduce the tasks required by a database administrator (DBA) in comparison to other relational database management systems. For example, with DB2 for i:

- Data Partitioning for performance/scalability: **NOT REQUIRED**
- Moving data or indexes to avoid disk hot-spots: **NOT REQUIRED**

- Re-balancing Indexes: **NOT REQUIRED**
- Monitoring table spaces, log buffers, lock contention buffers: **NOT REQUIRED**
- Running integrity checkers or statistics collection routines: **NOT REQUIRED**

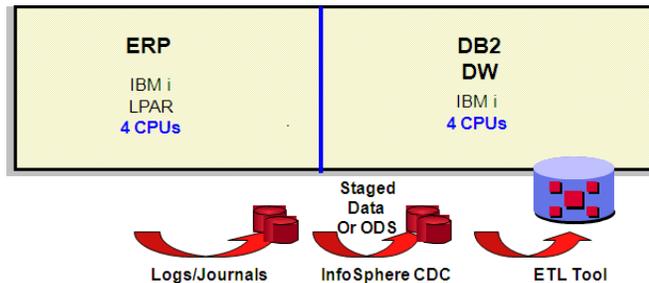
With IBM i navigator, the graphical administrative tool that comes with IBM i, capabilities exist to understand how DB2 is processing the BI workload. Real time analysis of query jobs through the SQL Plan Cache, or post execution analysis of SQL statements against collected database monitor data provide IT with graphical tools to understand how DB2 is processing requests by the BI application. Visual Explain, Index Advice, and additional analysis features make it very easy to tune the workload to improve performance and ensure your system is running optimally.

“The tools alone made the move from SQLServer to DB2 worth it. While SQLServer provides easy to use tools, the depth to really analyze what is happening with SQL based programs was very limited compared to DB2.” - Steve Hester, Senior Development Manager, Merillat Corp.

● **Virtualization Provides FLEXIBILITY**

Querying production systems may be the most affordable and realistic implementation option for many small and medium businesses. DB2 Web Query Standard Edition is a complete Business Intelligence package that installs easily into its own subsystem, providing real time analytics while adhering to DB2 for i governing settings to minimize impact of querying production database.

DB2 Web Query Standard Edition can also be implemented in a separate Logical Partition that can isolate the reporting server workloads from the production workloads. Leverage PowerVM technology to easily move processor or memory resources between the production and BI partitions to optimize processing based on dynamic workloads.



Completely isolate the BI workload from production with a separate server that can also be utilized for other purposes. The above mentioned IBM i for BI includes DB2 Web Query Standard Edition and data replication software and services to get an isolated operational data store implemented in a manner of days. Extend the IBM i for BI solution into an enterprise data warehouse by adding one of the several DB2 Web Query compatible ETL (Extract, Transformation and Load) tools available from IBM Business Partners.

● **OPTIMIZE your Investments**

Business Intelligence Applications can be a very powerful enabler for any company looking to gain insights into their data. Risk in any BI project is significantly reduced if you can leverage attributes common with current operational systems. These attributes include sharing of common hardware and software components, leveraging existing operational and administrative skills and policies, and being able to optimize investments across both operational and business intelligence applications through sharing resources and skills.

IBM i, Power Systems, and DB2 for i provide an excellent platform for minimizing risk and costs associated with implementation of this powerful technology. For additional information, please refer to:

DB2 Web Query for i Home Page: <http://www.ibm.com/systems/i/db2/webquery>.



All performance information was determined in a controlled environment. Actual results may vary. Performance information is provided "AS IS" and no warranties or guarantees are expressed or implied by IBM. Buyers should consult other sources of information, including system benchmarks, to evaluate the performance of a system they are considering buying.

Photographs show engineering and design models. Changes may be incorporated in production models.

Copying or downloading the images contained in this document is expressly prohibited without the written consent of IBM.

The Power Architecture and Power.org wordmarks and the Power and Power.org logos and related marks are trademarks and service marks licensed by Power.org.

UNIX is a registered trademark of The Open Group in the United States, other countries or both.

Microsoft, Windows, Windows NT and the Windows logo are registered trademarks of the Microsoft Corporation.

Intel, Itanium and Xeon are registered trademarks and MMX and Pentium are trademarks of Intel Corporation in the United States and/or other countries.

Java and all Java-based trademarks and logos are trademarks of Sun Microsystems, Inc. In the United States and/or other countries.

© IBM Corporation 2010
IBM Corporation
Systems and Technology Group
Route 100
Somers, New York 10589

Produced in the United States of America
March 2012
All Rights Reserved

This document was developed for products and/or services offered in the United States. IBM may not offer the products, features, or services discussed in this document in other countries.

The information may be subject to change without notice. Consult your local IBM business contact for information on the products, features and services available in your area.

All statements regarding IBM future directions and intent are subject to change or withdrawal without notice and represent goals and objectives only.

IBM, the IBM logo, ibm.com, DB2 are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. If these and other IBM trademarked terms are marked on their first occurrence in this information with a trademark symbol (® or ™), these symbols indicate U.S. registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of IBM trademarks is available on the Web at "Copyright and trademark information" at www.ibm.com/legal/copytrade.shtml

Other company, product, and service names may be trademarks or service marks of others.

IBM hardware products are manufactured from new parts, or new and used parts. In some cases, the hardware product may not be new and may have been previously installed. Regardless, our warranty terms apply.

This equipment is subject to FCC rules. It will comply with the appropriate FCC rules before final delivery to the buyer.

Information concerning non-IBM products was obtained from the suppliers of these products or other public sources. Questions on the capabilities of the non-IBM products should be addressed with those suppliers.

When referring to storage capacity, 1 TB equals total GB divided by 1000; accessible capacity may be less.

The IBM home page on the Internet can be found at: <http://www.ibm.com>.

The IBM Power Systems home page on the Internet can be found at: <http://www.ibm.com/systems/power/>