

IBM's System i5:

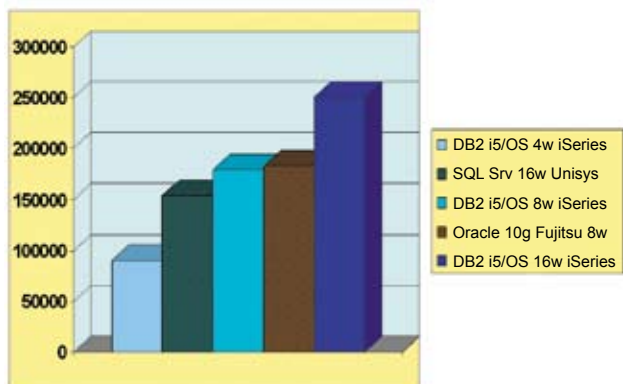
Key Attributes and Considerations For Business Intelligence Applications

Business Intelligence is a broad term relating to applications designed to gather, consolidate, cleanse and analyze data for purposes of understanding and acting on the key metrics that drive profitability in an enterprise. Examples include:

- Corporate Performance Measurement reporting, including Balanced Scorecard reporting and standard financial reports.
- Budget and planning applications that simplify the process and insure integrity of data during each planning cycle.
- Analytical CRM (Customer Relationship Management) applies BI practices to the marketing function to optimize budgets, execute and track campaign strategies, and generate more profit from marketing activities.
- Sales Analysis to track sales vs. forecast, to build plans based on profit or packaging strategies, and nurture premier customers with better service through a more informative sales force.
- Even Human Resource departments leverage BI applications to monitor and analyze payroll classifications, training programs or labor utilization.

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.

SAP BW 3.5 Benchmarks End User Query Performance



See detailed certified benchmark results at http://www.sap.com/solutions/benchmark/BW2_results.htm. Figure 1.

Key Benefits of System i for BI

DB2 for i5/OS Query Processing OUTPERFORMS!

The sophisticated cost based **query optimizer** built into DB2 for i5/OS 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 i5/OS’ cost-based optimizer uses statistics stored automatically in the database as well as other information, such as system configuration, parallelism settings, and available indexes in its development of the access plan.

“Pre-SQE it would have taken us 2.5 hours to process the high volume of data and generate the critical management reports. Our users expect us to miss our normal 7am delivery deadline on ‘Megadays.’ With the SQE code the processing completed in 55 minutes allowing us to meet the deadline much to the surprise and delight of our users.” - Andre Artymiuk, Database Administrator

IBM has been enhancing DB2 for i5/OS 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. An example of recent technology shipped with DB2 V5R3 (and enhanced in V5R4) is the new SQL optimization and statistics engine (SQE).

A great proof point of the performance capabilities of DB2 for i5/OS in a Business Intelligence application setting is the recent results of the SAP Business Warehouse benchmark. (See **Figure 1.**) The chart shows the consistent performance of the System i, over time, running SAP’s data warehouse solution, with the System i5 POWER5+ and DB2 for i5/OS V5R4 breaking the barrier in End User Query performance!

Another unique advantage of DB2 for i5/OS is IBM’s patented **Encoded Vector Indexing (EVI)** technology, built into DB2. EVIs use a symbol table and vector array indexing technology that can make complex data warehouse queries scream! For example, in a recent test in IBM’s DB2 Competency Center, the use of EVIs enabled a query of a 350 Gigabyte star schema database to run in less than 5 minutes compared to nearly 5 HOURS without EVIs! (See **Figure 2.**)

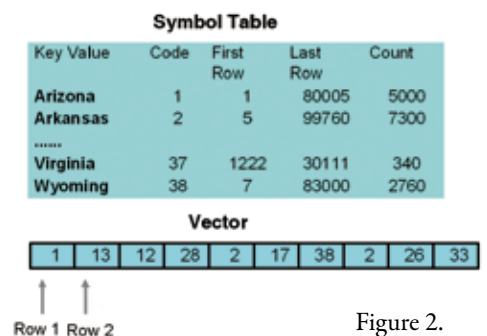


Figure 2.

Additional enablers for processing complex query workloads include:

- **Materialized Query Tables**
(MQTs)MQTs are “optimizer aware” summary tables. For large relational data warehouses, building of summary tables often speeds up commonly run queries. 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.
- **Remote Journaling**
“Near” real-time data warehousing can be accomplished through the use of a built in DB2 feature called Remote Journaling. Remote

Journaling allows capturing of changed database records real time into a log (Journal Receiver). Through i5/OS directives, the log can be shipped to a 2nd system or partition running i5/OS, and the change can be applied to the data warehouse based on timings that meet the requirements of the application.

• **Autonomic Indexes**

“Autonomic” functions revolve around the idea of the system being self-healing, self-managing, and self-tuning (as well as self protecting). DB2 for i5/OS has a long history of providing autonomic capabilities within the product to minimize Database Administration (DBA) tasks. Autonomic indexes are indexes that are built “on the fly” and preserved for future use — all done automatically by DB2.

DB2 for i5/OS lowers Total Cost of Ownership compared to other solutions

The System i5’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 i5/OS 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 i5/OS:

- 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**

“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.

The **On Demand Performance Center**, built into DB2 for i5/OS, provides capabilities to understand how DB2 is processing the BI workload. Real time analysis of the SQL Plan Cache, autonomic indexing, and easy to use Database Monitors help ensure your system is running optimally.

On Demand attributes of System i5 provide FLEXIBILITY

Logical Partitioning (LPAR) allows you to dynamically allocate resources (e.g., CPU or memory) across multiple copies of i5/OS (or Linux or AIX operating systems) within a single system footprint. The “on demand” features of LPAR also provide automatic resource allocation based on policies and priorities that you define.

LPAR technology can be applied to a Business Intelligence environment to reduce the Extract/Transformation/Load (ETL) and summary table/cube calculation times typically required in a data warehouse architecture. These workloads, typically running at night, can steal processor and memory from operational

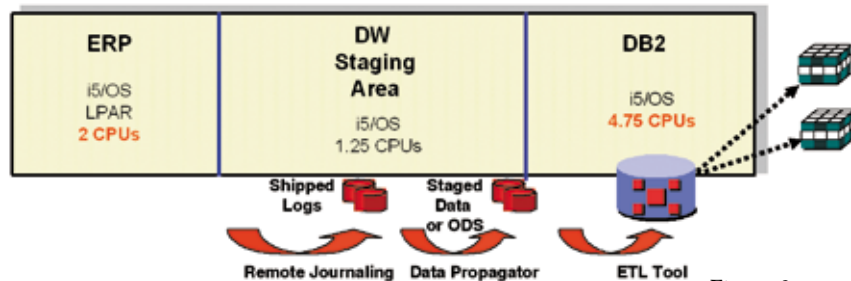


Figure 3.

application partitions that are going into batch or reduced workload modes. Once the ETL work is completed, those resources can be automatically moved back into the production partition for normal 8 to 5 workloads. **Figure 3** depicts how you could use Logical Partitioning in a BI environment.

Capacity Upgrade on Demand (CUoD) is a System i feature that allows you to quickly add additional processing power to your system, on a permanent or temporary basis. One of the challenges with BI applications is that usage of the application can grow like a snowball. Being able to react quickly to additional demands placed on the database server can be the difference between the success and failure of the application. With CUoD, you have almost instantaneous ability to turn on additional processors or acceleration features to meet these demands.

i3 Tech Group Inc.

Providing the *Right Answers* and the *Right People* to ensure your projects are on time and on budget

- Upgrades & Migrations - Hardware, Data and Application
 - Performance - Tuning, Reporting and Evaluations
 - Security Assessments - SkyView Risk Assessor
 - HMC and LPAR Implementation & Planning
 - Disaster Recovery Planning & Testing
 - BRMS Solutions - BRMS Partner
 - Linux & Windows Integration
 - Total Project Management
 - e-Business Solutions
 - Support Contracts
 - Education

IBM Certified Experts from V4R3 to V5R3
Canadian agent for CCSS Monitoring products

www.i3tg.com
905-841-2353

Scalability of DB2

A key consideration in any database platform in support of BI applications is its ability to scale while maintaining acceptable performance. DB2 for i5/OS offers a variety of techniques to address scalability/performance. **Parallel database** operations through the **DB2 Symmetric Multiprocessing (SMP) feature of i5/OS** 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 System i5 uses intelligent I/O Processors with their own levels of caching and parallelism. Combined with Single Level Store, the I/O subsystem of the System i removes a common bottleneck with data warehouse workloads: the I/O subsystem. The graph in **Figure 4** 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.

Combining High Availability or Disaster Recovery with Business Intelligence

Customers are often looking to enhance the availability of the System i5 based applications by leveraging a 2nd server and appropriate mirroring software to create a 24 x 7 or other disaster recovery scenario (in addition to being able to maintain availability during PLANNED downtime).

In addition, one of the challenges of the data warehouse architecture is how you TRANSPORT the data from source system to target while not impacting the production systems.

As with Remote Journaling (mentioned above), high availability software solutions for System i5 can also be leveraged to create an architecture that can provide near real time TRANSPORTATION of data from source to target systems, with minimal impact to the source side database server. An Operational Data Store (ODS), commonly used for real time operational reporting, is a BONUS of this architecture. The ODS can then also become the source for feeds into the data warehouse.

The benefits include:

- Data Transport becomes a non issue
- Leverage an otherwise UNDER utilized 2nd server for the data warehouse workload

- Achieve a Disaster Recovery or H/A solution for the data warehouse (as well as the operational applications)
- SAVE MONEY by combining the acquisition of a server for both H/A-DR and Business Intelligence

Leverage tools or application tiers running in Linux, AIX or Windows environments

Logical Partitioning (LPAR) described above, also supports Linux and AIX operating systems. For example, multiple i5/OS, Linux, or AIX partitions can be defined to a single System i5 server, reducing complexity, leveraging dynamic allocation of resources and other benefits such as virtual I/O, high speed interconnects between the partitions, and disk virtualization.

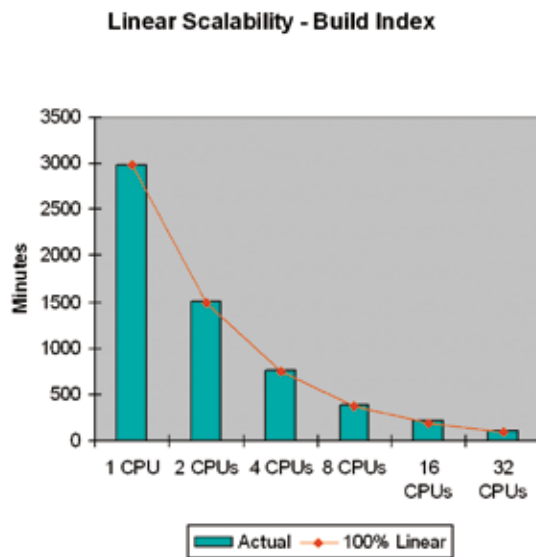


Figure 4.

virtual I/O and high speed connections to the server.

How does this relate to support of Business Intelligence applications? Suppose you are using an ETL tool whose TRANSFORMATION component does not support i5/OS but can run in Linux or AIX. Through support of multiple operating systems within a single System i5 server, the ETL code logic remains within the framework of the server. This implies that no data has to be moved across a network, and high speed bus communications can be leveraged to optimize the entire ETL process. All data – source databases, target databases, and any intermediate data storage areas, all remain within the framework of the System i5.

In addition, the use of the System i5’s integrated x-Series adapters can provide similar benefits to software designed for Intel processors by its ability to leverage the

“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 UDB for i5/OS was a better choice.” – Sam Gottlieb, Project Lead, Elie Tabari

Reduce Risk

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.

The IBM System i5 is an excellent platform for minimizing risk and costs associated with implementation of this powerful technology. For additional information, please refer to: DB2 for i5/OS Home Page: <http://www.ibm.com/servers/eserver/series/db2>. 