



White Paper

Sep 2008

Win in the flat world

Transforming Business Intelligence using SAP NetWeaver® Business Intelligence Accelerator and Intel® Xeon® Processors

As organizations juggle their business intelligence needs – some strategic, few tactical and others operational – they are hard pressed for options. Many organizations that run on SAP are unable to fully utilize the speed and flexibility of SAP Business Intelligence owing to their concerns about using it as an enterprise data warehouse. SAP NetWeaver Business Intelligence Accelerator* powered by Intel Xeon processors introduces much needed scalability and performance enhancement to the SAP BI ecosystem making it an ideal fit in many enterprise data warehouse environments.*

This paper discusses business scenarios where SAP BI Accelerator and Intel Xeon processors can help improve performance, and shares the performance related findings with and without SAP BI Accelerator using Intel Xeon processors.

For more information, Contact askus@infosys.com



Executive Summary

With Business Intelligence (BI) emerging as a critical success factor for strategic, tactical and operational information users, access to information alone is no longer enough. Organizations have been using BI for many years to monitor, report on, analyze, and improve the performance of their business operations. The focus of strategic technology leaders is to create solutions that can be integrated seamlessly across the IT landscape making them globally visible. In the flattening world, it is critical for companies to maximize application usage, better assimilate operational data, leverage real time data analytics for enhanced operational efficiency, and extract value from information for competitive advantage.

SAP's NetWeaver BI* suite helps enable organizations to create and operate a data warehouse in an enterprise-wide environment. Facilitating the integration of heterogeneous systems, it supports various system topologies and enables both strategic analysis and operational reporting.

In the past, SAP-BI implementations have often taken a module-centric approach with underlying challenges due to project-driven solutions, heterogeneous architectures, unmanaged business rule redundancies, uncontrolled data flow, and bottom-up/top-down approaches. The result was high development and maintenance costs. This created a compelling need for an integrated enterprise-wide data warehouse to ensure a 'single version of the truth' with assured scalability and efficient query response.

Current business demands require processing very large amounts of data to generate relevant analytical reports. Usually OLAP cache, RDBMS aggregates and pre-calculated queries are used to improve query performance. When data volumes increase and query navigation becomes sophisticated, it becomes challenging to provide consistently good query performance. Businesses are challenged by the need for large volumes of data, response time service level agreements and unpredictable query navigation.

Intel partnered with SAP in the development SAP NetWeaver Business Intelligence Accelerator* [BIA]* a technology powered by Intel Xeon processors that offers high performance and low maintenance for enterprises. This product is delivered as a packaged appliance, which is an optional component for SAP NetWeaver BI.

The BIA Appliance can benefit businesses that have high volumes of data. The read performance of queries can be greatly improved using the BIA Appliance. This new tool helps enable quick access to any data with a low amount of administrative effort and is especially useful for sophisticated scenarios with unpredictable query types, high data volumes and high frequency of queries. BIA is also useful when aggregates or database indexes are not sufficient, or when these methods become too costly to maintain.

As shown in this paper, the advantages of the BIA Appliance include:

- Consistent high performing queries
- High scalability

The performance gain using the Intel/SAP BIA Appliance can be significant, potentially boosting the query response times by a factor of 10 to 100 times.

The BIA Appliance provide performance enhancement in many scenarios but not all. BIA does not support real-time reporting and, in the current version, is fetching data only from basic InfoCubes and not from virtual or remote InfoCubes, data store objects and source systems.

Need for Intel/SAP BIA Appliance

Below are the key business reasons for using the BIA Appliance:

1. Growth in number of end-users

- Internal access expanding beyond specific analyst user community
- About 44% of companies provide BI reports to their customers, suppliers, stakeholders etc.¹

2. Growth in data volumes

- 30% of companies (revenue of \$500M or more) expect 100% growth in their data warehouses.¹
- Turning high-volume data to high-volume information

3. Shorter decision cycles requiring more flexibility

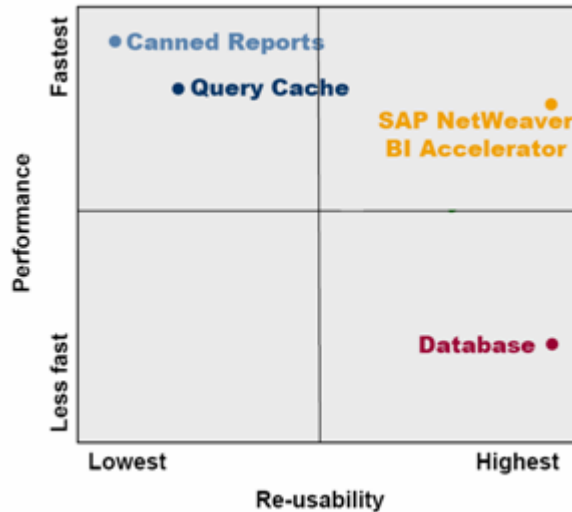
- More flexible ad-hoc querying support needed for better decision making
- Availability of actual data for forecasting

Blade systems from IBM* or HP* using Intel Xeon processors form the foundation for the BIA Appliance. The BIA Appliance provides exceptional performance due to the factors listed:

- Runtime in-memory aggregation
- Compact index structures stored in a separate file server
- Vertical decomposition (similar to column-oriented databases)
- Horizontal partitioning of indexes for high parallel processing on multiple processors
- Optimized for 64-bit OS
- Smart compression
- Proven TREX technology extended for structured data
- Blade server technology allows for dynamic resource allocation for high availability and load balancing

BIA Appliance Performance: Speed vs. Reusability

The BIA Appliance offers high reusability together with fast performance in a manner not possible with other techniques such as canned reports and query cache.



Source: see Endnote 2

Intel's Hardware Infrastructure

Intel's Extended Memory 64 Technology adds 64-bit capability for the Intel Xeon processors providing higher performance and expanded memory required by the BIA Appliance in a complete package.

Intel collaborated with SAP in the development of the BIA Appliance by using advanced Intel software and hardware technology optimizations to maximize performance and scalability in solutions based on enterprise services architecture on innovative 64-bit Intel platforms. Intel engineers provided 64-bit porting assistance for the BIA and extensive performance optimizations of the product. With the current multi-core technologies incorporated into the Intel Xeon processor family, the BIA Appliance provides even better performance and throughput.

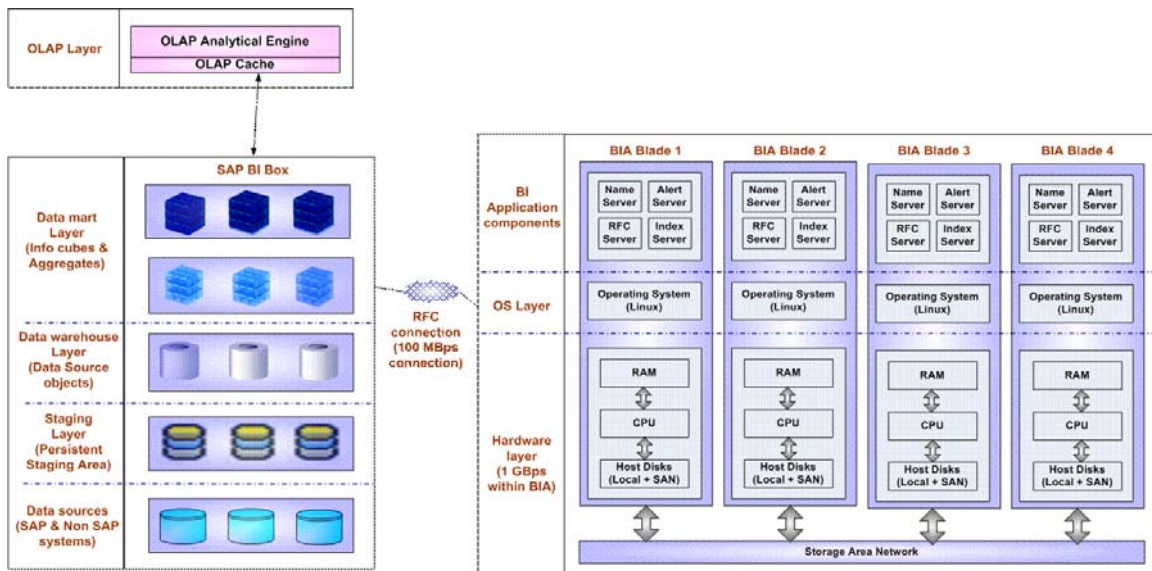
The BIA Appliance runs on Linux-based systems that employ high performance, 64-bit Intel processors and modular blade technology. These hardware platforms are currently available from IBM and HP:

- IBM Blade Center from IBM TotalStorage*
- HP Blade System from HP StorageWorks*.

Solution Architecture

The BIA Appliance, developed by Intel and SAP is a plug-in appliance for SAP Netweaver BI. It enhances the query performance manifold and improves the scalability of the SAP BI environment. The BIA Appliance is an Intel Xeon processor-based high-density blade system from HP or IBM, with pre-installed SAP software. The blades contain Intel processors and run SUSE Linux operating system.

Solution Architecture Diagram: BI Accelerator



The BIA Appliance is an installed TREX* system with a BIA engine. The data of the BI InfoCube is stored and processed entirely in the main memory of the BIA Appliance server. The BIA engine is the part of the analytics engine that manages the BIA index. This software allows the system to read data from the BIA index, add data to the BIA index, or change data. The BIA optimizer is the part of the BIA engine that ensures the best possible read access to a BIA index.

Normally, optimization of query response time requires some form of pre-aggregation. The BIA Appliance, however, performs aggregation at query runtime. With the BIA Appliance, all query processing is handled in memory.

The BIA Appliance first indexes each InfoCube to be accelerated. This index structure is highly compressed and loaded into memory whenever the data is requested by the query. The BIA Appliance processes any incoming queries against indexes in memory using fast join and aggregation algorithms based on efficient integer coding techniques. The results are then passed to the BIA Appliance analytic engine and returned to the user. To create the indexes, database tables are decomposed vertically and stored by column, not row-based, and the indexes can be split horizontally for parallel processing on multiple processors.

The blade server infrastructure allows for dynamic assignment of hardware resources to ensure high availability and good load balancing. The BIA Appliance engine holds table data vertically, into columns that are stored separately. This makes more efficient use of memory space compared to row-based storage, since the engine needs to load only the data for relevant attributes or characteristics into memory. This is particularly beneficial for analytics, where most users want to see only a selected amount of data.

In a conventional database, all the data in the table are loaded together, in complete rows. The BIA Appliance engine touches only relevant data columns. The engine can sort the columns individually to bring specific entries to the top. The columns indexes are written to memory and cached as flat files. Efficiency is improved because both the memory footprint of the data and the input-output flows are smaller.

Vertically segmenting the data tables in this way is more efficient than saving row-based data in conventional relational database systems. In a conventional database, the system has to search all the data in the table.

The BIA Appliance engine can partition large tables, such as Fact tables, horizontally for parallel processing on multiple machines in distributed landscapes. This enables processing of huge data volumes within the limits of installed memory. Volumes of data can be split over multiple hosts, by a round-robin procedure to build up parts of equal size, so that they can be processed fast and in parallel. A logical index server distributes join operations over partial indexes and merges partial results.

BIA data is stored in a read-optimized format. The data is compressed using integer coding and dictionary lookup. Integers represent the text or the values in the table cells, and the dictionaries are used to replace integers by their values during post-processing.

In particular, each record in a table has a document ID, and each value of a characteristic or an attribute in a record has a value ID. An index for an attribute is simply an ordered list of value IDs paired with sets of IDs for the records containing the corresponding values. Integer compression greatly reduces the average volumes of processes and cached data. This allows more efficient numerical processing and smart caching strategies, which reduce the data volumes and flows by an average factor of ten. Hence all processing can be performed in the main memory.

Business Scenarios

Some business scenarios where the Intel/SAP BIA Appliance could be a critical success factor:

High data volumes

Almost all companies in retail and utilities sectors store enormous amount of customer data so that they can run analytics tools on them so as to gather intelligent insights. Since this data runs into millions of rows or more, performing analysis on this data becomes increasingly difficult. Hence the BIA Appliance is extremely useful in such scenarios to carry out analysis on the data.

Challenging response times for Service level agreements

Service level agreements for industries that require extremely small turn around times so that customer requests can be effectively closed thus making the BIA Appliance ideal for companies working in such sectors.

Unpredictable type of queries

The BIA Appliance is also useful in situations where we come across requests for on-demand reporting from user groups by ad hoc analysis that require different data sets and aggregations than traditional optimization and caching strategies. The BIA Appliance helps in reducing the response times for drill-down, slice and dice in such scenarios.

Business Scenarios Implemented in the POC

Demand Change Scenario

Demand planners at any site need to plan for material demand across a period of time. The data mart typically contains data from multiple sources such as planned orders, purchase orders, purchase requisitions, etc. For material in a given plant it is calculated as 'old/new demand' and the corresponding change for a particular item 'sold to party'. Execution time for a typical query to calculate the change in demand for a period of a month would be 5 minutes for a cube with 500 million records.

Sales Trend Analysis and Drilldown

This analysis shows 'sales quantity' and 'sales value' over a period of time. The requirement is to be able to view this data through various dimensions like geography, product group, POS, etc., and drill down to various levels on these dimensions.

The BIA Appliance is loaded with high-performance software in the SAP NetWeaver platform with advanced hardware from Intel to deliver performance boosts of ten or more times for analytic services and analytic applications.

Demand Cube

The demand cube stores the demand data for all the materials for a period of six months to the current date. The demand data is derived from Customer Orders, Planned Orders, Purchase Orders and Purchase requisitions.

| Demand cube details | |
|---------------------------------------|------------------------------------|
| 70 Million Rows Size: 9,167,912 KB | |
| Dimensions of the Demand Cube | Key Figures for Demand Cube |
| Material | Order Quantity |
| Plant | Inherent High |
| Customer | Event 60% Positive |
| Sales Document | Budget |
| Additional Time characteristics | Current Sales Plan |
| Sales Area | Inventory |
| Material, plant view | Dependant Order Quantity |

Sales cube

The sales cube contains sales of material over a time period. The user can drill down on various dimensions like geography, product group, POS etc and also drill down to various levels on these dimensions.

| Sales cube details | |
|--|--|
| 94 Million Rows Size: 20,223,504 KB | |
| Dimensions of the Sales Analysis Cube | Key Figures for Sales Analysis Cube |
| Material | Quantity in base units of measure |
| Customer | Net Value |
| Sales Document Category | |
| Additional Time characteristics | |
| Unit | |
| Organization | |
| Document Classification | |

The Sales and Demand Cube consists of huge data as they store the demand for all the materials. They store the data on a daily basis and the data just gets added on every day and the number of records runs into millions.

This makes the execution of simple looking reports very hard. The reports take a lot of time to execute and the user has to wait almost 10-15 minutes for the output. This accounts for poor usability for the end-user. Hence BIA Appliance is needed for a better user experience.

Current SAP BIA Set-up

a. System information

The current set-up consists of a 4 blade (each with 16GB RAM and 144GB hard disk) system with a 1TB SAN.

Details:

2x3.0 GHz Dual Core Intel Xeon 5160 Processor

8x2 GB PC2-5300 CL5 ECC 667MHz FB DDR2 Chipkill Memory

146 GB 10K U320 FSS SAS Hard Drive

2x Broadcom Gigabit Ethernet

Brocade 10-port 4Gb SAN Switch Module

QLogic 4Gb Fibre Channel Expansion Card

IBM System Storage DS4700 Express Model 72

146.8 GB 15Krpm 4GB/s FC E-DDM Option

BIA: Suse Linux Enterprise Edition 9

BI: Microsoft Windows server 2003

SAP BIA Version: 700.36.119243 where BIA release: 700, BIA revision: 36, internal version number: 119243

b. Sizing

To ensure optimal performance, sizing is done prior to the purchase of the BIA Appliance as it holds indexes. When the number of cubes (data targets) for which the indexes are created changes, we might need to accordingly resize the BIA Appliance by adding additional blades/memory).

If both sales and demand cubes reside together in the BIA Appliance then it approximately needs 20GB of space.

Below are the sizing recommendations

| Value | Size |
|---------------------------|------------------------------------|
| Required Memory Value | 20 GB |
| Total Memory Value | 2 * 20 GB = 40 GB |
| Number of Blades required | 40 GB / 16 GB per blade = 3 blades |
| Total Storage Value | 3 * 20 GB = 60 GB |

Since the current set up of BIA Appliance has a 4 Blade system with 64 GB RAM, it is sufficient for handling both the sales and demand cube.

Note: *Memory consumption can be calculated using the ZZ_BIAMEMCONSUMPTION_BW3X report in SAP. The report option P_DETAIL gives more precise results related to sizing.*

c. Connectivity

Currently there exists 100 Mbps link between the SAP BI system and BIA Appliance. Even though SAP recommends a 1Gbps network link between these two systems, the current infrastructure is sufficient since real time updates of data is not occurring. All the queries are run against static data loaded once.

Learnings from BIA Implementation

The objective of the effort undertaken at the BIA lab, Infosys was mainly to test and compare query performances of BIA in different scenarios. In this process the team studied BIA at a detailed level and developed competency in the following

- Understanding the likely scenarios where BIA is most suited for optimal query performance.
- Interfacing of the BIA Engine with internal, as well as other external, systems.
- Understanding the pre requisites for optimal usage of BIA.
- Resolving operational bottlenecks while using BIA.
- Understanding the impact of hardware/configuration changes in BIA on Query performance.

Engineers at Infosys consolidated the test results as well as lessons learnt and best practices and delivered the following

- BIA Implementation Methodology and Test Scenarios
- Pre-requisites and Operational & Configuration guidelines for BIA
- Best practices for usage of BIA.

Results

During the testing a series of queries were run at the Infosys SAP lab, Bangalore, in order to study the impact of BIA Appliance. Tests were performed to analyze the impact of BIA Appliance on both stand-alone queries and on multi-user scalability.

The major items that contribute towards total query response time in BIA are Data manager time (DM Time), OLAP time and Front end time.

| Times | Event Text |
|--------------------------|---|
| OLAP Time | Cache Generation |
| | Cache Commit |
| | OLAP: Settings |
| | OLAP: Query Gen. |
| | OLAP: Read Data |
| | OLAP: Data Selection |
| | OLAP: Data Transfer |
| | OLAP Initialization |
| | OLAP: Read Texts |
| | OLAP Other Time |
| Data manager time | Data Manager |
| Front end time | 3.x Analyzer Server |
| | 3.x Query View Open |
| | 3.x Query View Data |
| Others | Total DBTRANS (Total Number of Transported Records) |
| | Total DBSEL (Total Number of Read Records) |
| | RFC call |

Stand-alone Queries

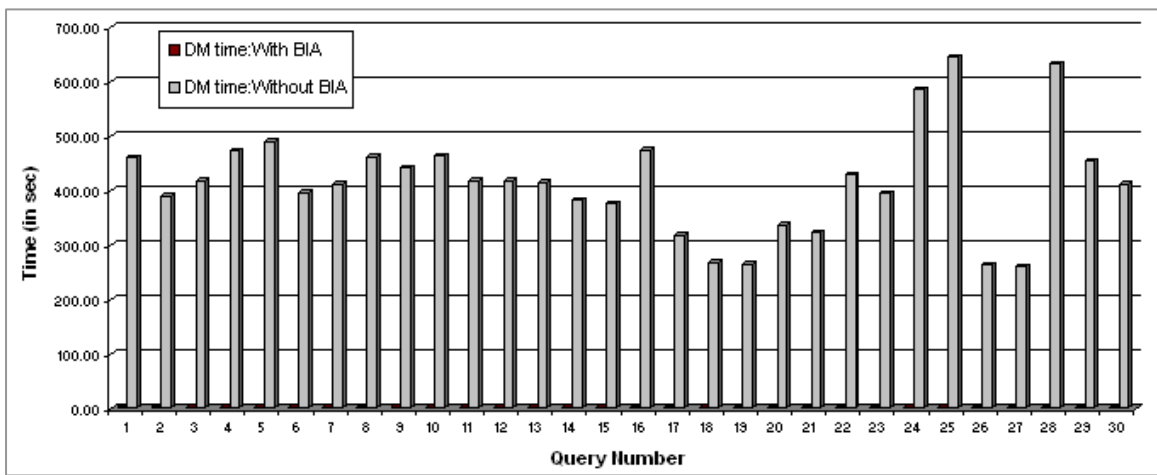
A set of 30 diverse queries based on sales and demand InfoCubes were chosen and query response times were noted. The total response times were noted for each query with and without the BIA Appliance. In each and every case the queries run with BIA Appliance outperformed the same queries run without the BIA Appliance.



The total query response time with BIA Appliance was on an average 205 times faster than the total query response time without the appliance.


a. DM Time

The data manager time is the time taken to read the data from BIA Appliance. The data manager time is most relevant while comparing results with and without BIA Appliance. From the results we can observe that for every query that was run with BIA Appliance the DM time was found to be much better than DM time without the appliance. On average, a 396 times improvement was observed in DM time while using the BIA Appliance.



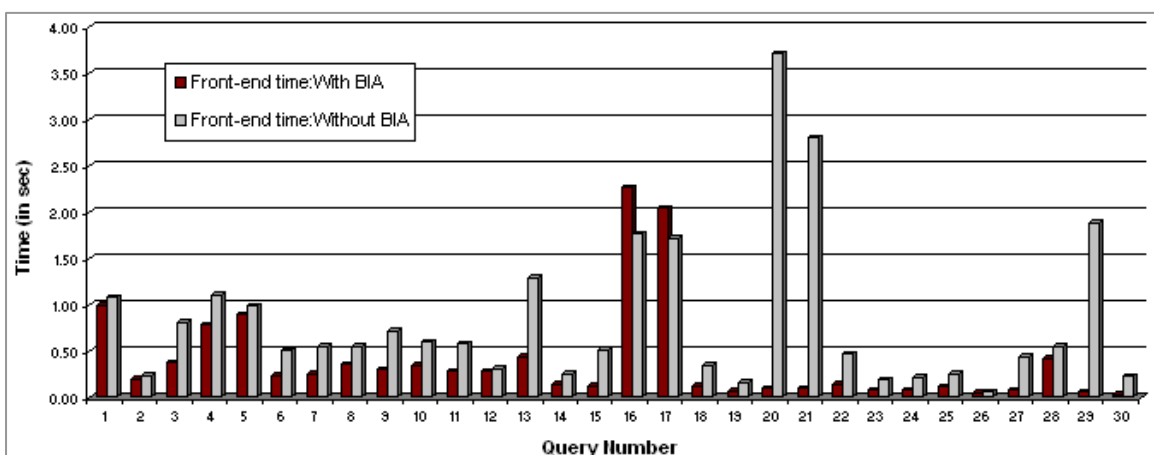
b. OLAP Time

OLAP data transfer time contributes the most towards the OLAP time calculation. During OLAP data transfer time exception aggregations and simple currency translations are carried out, formulas are calculated, and the correct number of decimal places for the data cells is determined. In addition, the result set is sorted according to the interface settings. In 94% of cases (28 out of 30 queries) the OLAP time with BIA Appliance was better than OLAP time without BIA Appliance. The average OLAP time for queries with BIA Appliance was found to be 3.4 times better than OLAP time without BIA Appliance.



c. Frontend Time

Front end time counts the time to transfer data to the front-end and the time to format the data in the front-end. Front end time consists of 3.x analyzer server, 3.x query view open and 3.x query view data events. 3.x analyzer server event measures the time for the logic at the ABAP server. 3.x query view open event measures the time taken to generate a data provider in the 3.x runtime. 3.x query view data event measures the time taken to read the data in the 3.x runtime. The front end time with BIA Appliance was found to be almost the same as or better than that of front end time without BIA Appliance. The average FE time for queries with BIA Appliance was found to be 5.4 times better than FE time without BIA Appliance.



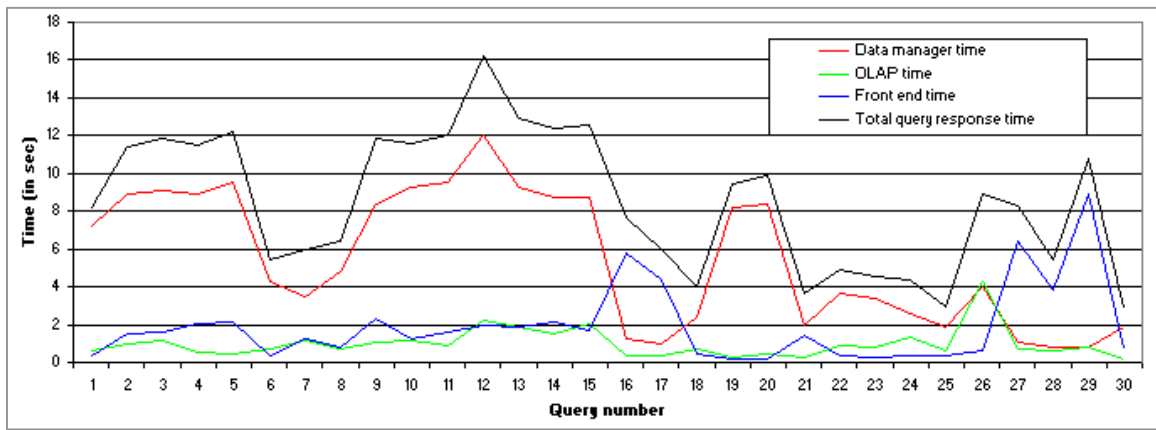
Concurrent User Queries

The same set of 30 queries was run in concurrent mode (10 users) to measure the scalability that BIA Appliance could handle. The LoadRunner* test utility was used to test the performance of the appliance in concurrent mode.


Loadrunner settings: The load runner was set to record mode and the script for load runner was recorded. Recorded script was then edited manually to parameterize variables like user name, password, query names (for sequential execution), export file name and path in which the exported files needed to be stored etc. The load runner controller was used for running the script for multiple users. In this step, run time setting and concurrent user initialization was performed.

Multi user mode: The average data manager time, OLAP time, Front end time and Total query response time in multi user mode were as follows:

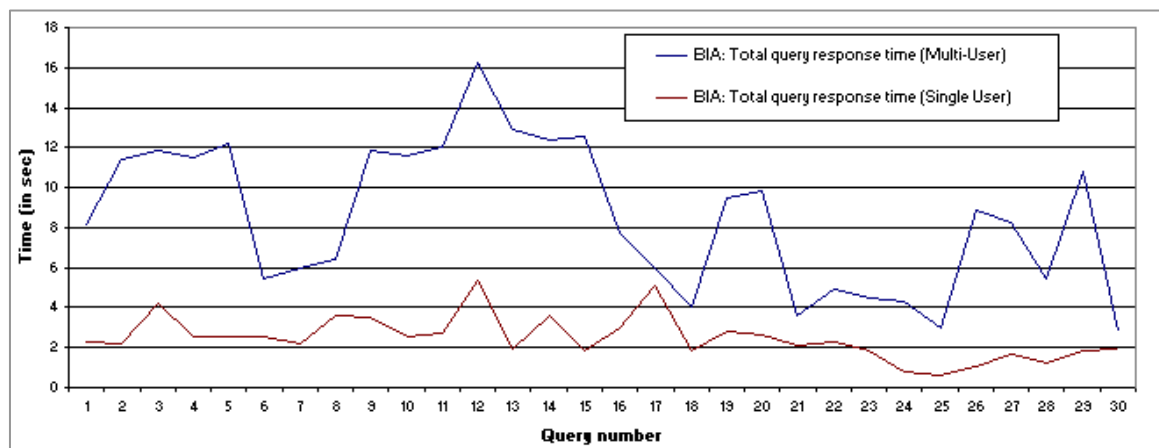
| | |
|-----------------------------------|-----------|
| Average DM time | 5.51 secs |
| Average OLAP time | 0.99 secs |
| Average FE time | 1.9 secs |
| Average Total query response time | 8.53 secs |



The performance of queries with BIA Appliance in the multi user mode improved by a factor of up to 54.25 times and the average query response time was 481.49 secs.



As expected, the total query response time with BIA Appliance in single user mode was better than total query response time without BIA Appliance in multi user mode. The average query response time with BIA Appliance in single user mode was 2.46 secs where as the average query response time with BIA Appliance in multi user mode came out to be 8.35 secs.

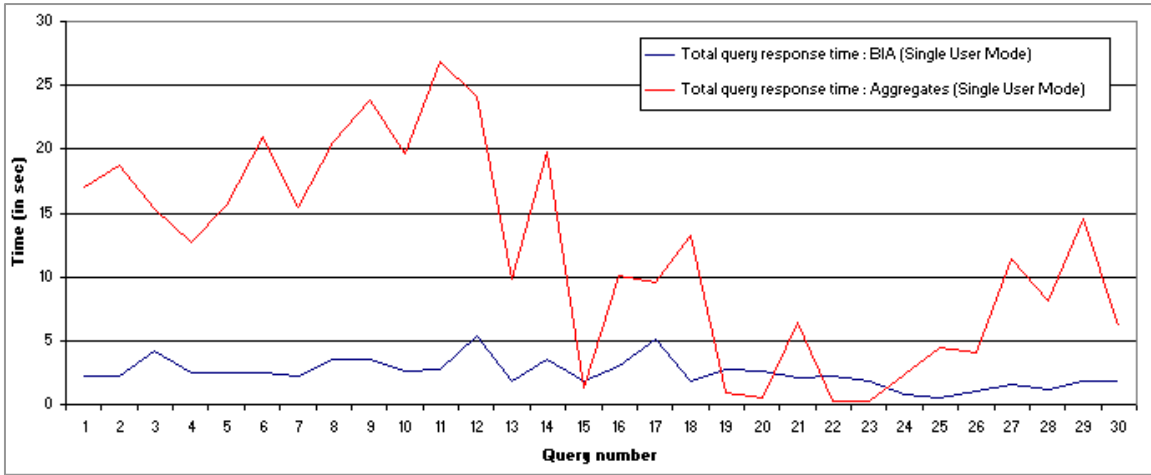


a. Aggregates vs. BIA

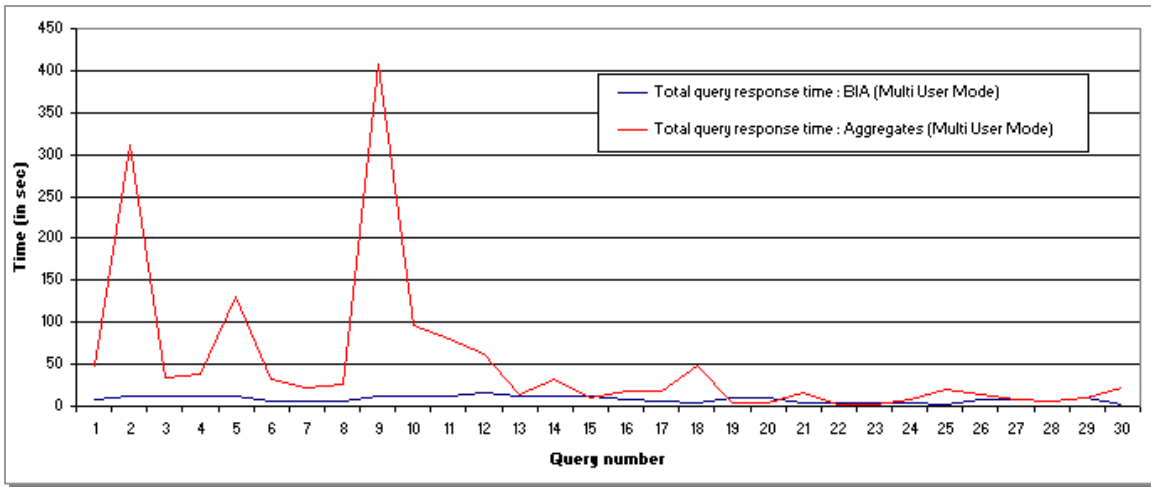
All test queries were run with both aggregates and with BIA Appliance. BIA Appliance performed much better than aggregates in

- 84% of cases (25 out of 30 queries) in single user mode and
- 77% of cases (23 out of 30 queries) in multi user mode.

The average total query response time taken by aggregates in single user mode was 11.814 secs where as the average total query response time taken by BIA Appliance in single user mode was 2.46 secs (improvement factor of 4.9 times).



The average total query response time taken by aggregates in multi user mode was 51.63 secs whereas the average total query response time taken by BIA Appliance in multi user mode was 8.53 secs (improvement factor of 5.5 times). For all those cases in which aggregates outperformed BIA Appliance, they could do so only by a very small factor.



Aggregates tend to outperform BIA Appliance in cases where an earlier version BIA application is used. The BIA application revision number for the SAP BIA implemented at Infosys SAP labs is 36. The performance of BIA with respect to aggregates is expected to be much better with later revision numbers of BIA.

Infosys Recommendations for Intel Xeon Processor based BIA Appliance

| | |
|--|---|
| <p><i>BIA Appliance is useful if it satisfies the following criteria</i></p> <p>Huge number of scanned records (i.e. DBSEL > 1 Million records)</p> <p>or</p> <p>High Aggregation ratio > 100</p> <p>or</p> <p>Large number of transferred records (DBTRANS > 50000)</p> | <p>DBSEL event measures the number of read records, aggregated from all the read accesses.</p> <p>DBTRANS event measures the number of transported records, aggregated from all read accesses.</p> <p>Aggregation ratio = DBSEL / DBTRANS</p> |
|--|---|

The BIA Appliance could be compared against aggregates on various factors such as query performance, ease of implementation, technical knowledge and ease of maintenance. BIA Appliance can be easily implemented as compared to aggregates since it is just a device that needs to be plugged in the current SAP BI landscape. Aggregates, on the other hand, need to be created exclusively on the InfoCubes for faster retrieval. Aggregates are likely to perform better than BIA Appliance if earlier revision numbers of the BIA application is used. For BIA Appliance using the application with revision number 40 or more the query performance is likely to be superior as compared to aggregates. The administrative effort required for maintenance of aggregates is much more as compared to BIA Appliance. The technical know-how of customization and configuration settings for BIA Appliance needs special attention and the implementation team needs fair amount of expertise for successful implementation.

Based on our experience with the BIA Appliance implementation at the Infosys lab, below table summarizes the evaluation of BIA Appliance vs. aggregates.

| Comparison parameters | BIA Appliance | Aggregates |
|---------------------------------|---------------|------------|
| Ease of Implementation | Fair | Good |
| Average query performance | Excellent | Good |
| Ease of maintenance | Good | Low |
| Technical knowledge requirement | Fair | Fair |
| Data Loading performance | Good | Fair |

The table below summarizes the performance comparison for one of Infosys' BI clients using the BIA.

| SAP BW | | |
|----------------------------------|-------------|-----------------|
| Total Size | BIA Server | 150 GB |
| Largest Cubes | Sales Cube | 94 million rows |
| | Demand Cube | 70 million rows |
| Response Time Without Aggregates | Sales Cube | 406.47 secs |
| | Demand Cube | 407.55 secs |
| Response Time With Aggregates | Sales Cube | 17.37 secs |
| | Demand Cube | 6.16 secs |
| Response Time With BIA | Sales Cube | 2.78 secs |
| | Demand Cube | 2.03 secs |

Customer Experiences

a. Major Agro based company -

Europe's leading agri-based client having a strong worldwide market presence in crop protection. Third in high-value commercial seeds .Excellent complementary and leading product portfolio having first class research and development resources and a rich pipeline of biotechnology traits in Plant Science.

Some of the challenges faced by the client included - Project driven solutions, Heterogeneous architectures, concepts & technologies, unmanaged redundancies of business rules, uncontrolled dataflows and bottom up/top down approaches and expensive developments and maintenance.

Infosys partnered with the client on the BIA Implementation and helped the client achieve significant improvements in query response times and also helped them put in a process in place for the BIA Governance process.

b. Leading UK based Retailer -

With around 2,600 stores from local community pharmacies to large destination health and beauty stores, this client is a world-leader in healthcare products and services.

Usage of the BI system was predicted to grow by 50% from Christmas 2007 to Christmas 2008 and it was proposed to upgrade BW to SAP NetWeaver Business Intelligence (BI), re-platform the BW hardware and implement the SAP BI Accelerator (BIA). These changes meant that the client has a scalable, fully-supported business intelligence solution that can meet predicted usage and volume demands, and provide more consistent and faster query performance.

Infosys partnered with the client on the BIA Implementation and support. Adding more infocubes data in scope for BIA, monitoring and analyzing reports out of BIA and interfacing with hardware vendor(s) on the BIA were some of the key activities performed.

Conclusion

The SAP BIA Appliance built upon high performance Intel Xeon processor technology is a very powerful appliance and aids customers who have problems with performance of queries in a SAP BI environment. Its ability to seamlessly integrate without any customization requirements of the existing system makes it the right choice for many customers. SAP BIA Appliance introduces much needed scalability and performance enhancement to the SAP Business Intelligence system.

Performance metrics from the experiments indicate that response times with BIA Appliance using Intel Xeon processors are far better when compared to same without BIA Appliance both in single user mode and concurrent mode.

The following points should be kept in mind regarding query performance related to BIA

- Query processing time basically comprises three components: data access including aggregation, calculations (OLAP), client rendering; BIA Appliance addresses the data access time. In some special situations it also addresses the calculation time
- Slow queries (when running without BIA) may significantly benefit by using the BIA Appliance; fast queries won't benefit that much. If the query is already fast without BIA Appliance then a significant improvement may not be experienced.
- Beside better performance, you may experience stable performance; this results from the simple processing paradigms inside BIA Appliance that
 - (a) Avoid the volatility of a DB optimizer and
 - (b) Achieves linear scalability: more blades or faster CPUs typically directly translate into better performance.

Summary of performance related findings are as follows:

- The query response time when queries were run with SAP BIA Appliance improved by a factor of up to 205 times more than when they were run without the SAP BIA Appliance (single user mode)
- BIA Appliance outperformed aggregates in up to 77% of test cases in concurrent mode. For the rest, aggregates outperformed BIA Appliance only marginally. The average total query response time taken by aggregates in multi user mode was 51.63 secs whereas the average total query response time taken by BIA Appliance in multi user mode was 8.53 secs (improvement factor of 5.5 times).

References

1. IDC Insight: SAP Business Intelligence Accelerator: A High-Performance Analytic Engine for SAP NetWeaver Business Intelligence. Author: Dan Vesset, [October 2005, IDC #34234, Volume: 1]
2. [SAP AG: SAP NetWeaver Business Intelligence Accelerator](#). Author: Alexander Peter, Product Management SAP NetWeaver BI, SAP AG. [2005, SAP Business Intelligence and Analytics Conference]
3. Winter Corporation: The SAP NetWeaver BI Accelerator – Transforming Business Intelligence. Authors: Rick Burns and Robert Dorin [September 2006]

About the Author

Meena Iyer is a Senior Consultant with Infosys' SAP Business Intelligence division of the Enterprise Services group. With experience in multiple geographies across several areas on SAP BI landscapes, she is a technical expert on SAP BW. Her team of BW and Basis talent worked on the proof-of-concept for 4 months to arrive at the detailed white-paper. Meena can be reached at meena_iyer@infosys.com



2008 Infosys Technologies Limited / Intel Corporation. All rights reserved.

Copyright in whole and in part of this document "**Transforming Business Intelligence using SAP NetWeaver® Business Intelligence Accelerator and Intel® Xeon® Processors**" belongs to Infosys Technologies Limited and Intel Corporation. This work may not be used, sold, transferred, adapted, abridged, copied or reproduced in whole or in part in any manner or form or in any media without the prior written consent of Infosys Technologies Limited and Intel Corporation.

Performance tests and ratings are measured using specific computer systems and/or components and reflect the approximate performance of Intel products as measured by those tests. Any difference in system hardware or software design or configuration may affect actual performance. Buyers should consult other sources of information to evaluate the performance of systems or components they are considering purchasing. For more information on performance tests and on the performance of Intel products, visit Intel Performance Benchmark Limitations

Intel processor numbers are not a measure of performance. Processor numbers differentiate features within each processor family, not across different processor families. See http://www.intel.com/products/processor_number for details.

Intel does not control or audit the design or implementation of third party benchmarks or Web sites referenced in this document. Intel encourages all of its customers to visit the referenced Web sites or others where similar performance benchmarks are reported and confirm whether the referenced benchmarks are accurate and reflect performance of systems available for purchase.

INFORMATION IN THIS DOCUMENT IS PROVIDED IN CONNECTION WITH INTEL PRODUCTS. NO LICENSE, EXPRESS OR IMPLIED, BY ESTOPPEL OR OTHERWISE, TO ANY INTELLECTUAL PROPERTY RIGHTS IS GRANTED BY THIS DOCUMENT. EXCEPT AS PROVIDED IN INTEL'S TERMS AND CONDITIONS OF SALE FOR SUCH PRODUCTS, INTEL ASSUMES NO LIABILITY WHATSOEVER AND INTEL DISCLAIMS ANY EXPRESS OR IMPLIED WARRANTY, RELATING TO SALE AND/OR USE OF INTEL PRODUCTS INCLUDING LIABILITY OR WARRANTIES RELATING TO FITNESS FOR A PARTICULAR PURPOSE, MERCHANTABILITY, OR INFRINGEMENT OF ANY PATENT, COPYRIGHT OR OTHER INTELLECTUAL PROPERTY RIGHT.

UNLESS OTHERWISE AGREED IN WRITING BY INTEL, THE INTEL PRODUCTS ARE NOT DESIGNED NOR INTENDED FOR ANY APPLICATION IN WHICH THE FAILURE OF THE INTEL PRODUCT COULD CREATE A SITUATION WHERE PERSONAL INJURY OR DEATH MAY OCCUR.

Intel, Intel logo, and Xeon are trademarks of Intel Corporation in the U.S. and other countries.

* Other names and brands may be claimed as the property of others.