

Optimized Performance with SAP Net Weaver BI Accelerator

M. Victoria Hebseeba¹, Dr. T. Bhaskara Reddy².

¹Department of Computer Science, Rayalaseema University, Andra Pradesh

victoria.hebseeba@gmail.com

²Department of Computer Science, S.K. University, Andra Pradesh

bhaskareddy_sku@yahoo.co.in

ABSTRACT

This paper describes the SAP Net weaver BI accelerator and its features and how the accelerator can be used for optimized performance in any IT industry. The objective is to provide the best practices and solutions for current challenges in using SAP Net weaver BI Accelerator in any large scale business industries across the world.

I. INTRODUCTION

SAP Net weaver BI Accelerator extends the traditional environment, crunching through terabytes of data in seconds, enabling faster business insight and turbo-charging your business intelligence solution. It allows quicker access to data that has been warehoused for varying periods of time which, in the past, has been too complicated to retrieve and utilize effectively.

The SAP Net weaver Business Intelligence Accelerator is an appliance combined in one package software and hardware. BI accelerator was developed by SAP in collaboration with Intel, which provides the processors, and HP and IBM, which provide their respective server and storage technologies. Appliance delivers enhanced business intelligence functionality while connecting to BI system landscape.

SAP BI Accelerator reduces query response time particularly for large data volumes by taking rows of data and placing them into modern columnar layout. BI Accelerator provides radical improvement in performance to both frequently and infrequently asked Queries in the SAP BI platform [1]. Compared with the other approaches the appliance offers a great improvement in speed and flexibility of data access offering a number of benefits to IT industries in terms of the increase in productivity as well as lowering the total cost of ownership and profitability.

The data in BIA indexes is consistent with the Info Cube data in the BI system. This can be assured by BI Accelerator Data Consistency Check from the BIA Monitor [2]. Prior to deployment it is important to know what you are delivering and what the BWA is competing against the solution.

If the basic BW model itself is defective and fundamental architecture is having flaws then BWA will provide lesser value. Most of the reports should be validated to ensure they are business friendly; else BIA will not improve data or information quality. This paper explains how the performance can be optimized by using the appliance SAP Net weaver BI Accelerator.

II. BI SYSTEM AND BI ACCELERATOR

SAP Net Weaver BI customers adopting the BI Accelerator can expect significant improvements in query performance through in-memory data compression and horizontal and vertical data partitioning, with no administrative overhead. BI Accelerator is delivered to the customer as a preinstalled and preconfigured system on dedicated hardware as a BI Accelerator box, therefore the installation and initial configuration has been done and no additional administrative tasks need to be done by the customer for the first usage of the BI Accelerator.

To ensure the uncompromised performance of BIA there is 1:1 connection between the BI instance and the BI Accelerator. BI Accelerator works as follows after connecting with BI Instance through Remote Function Call (RFC).

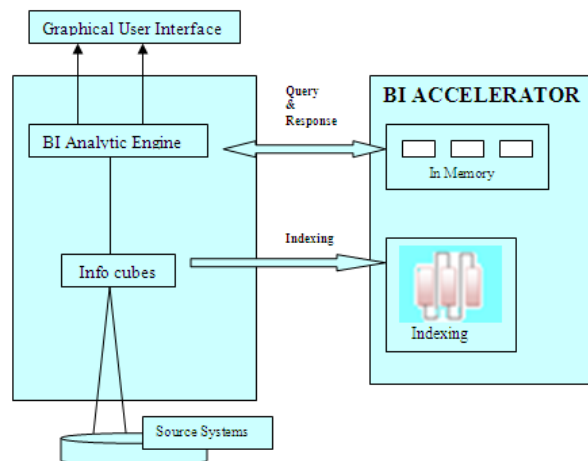


Figure1. BIA Architecture

- Data from an SAP Info Cube is loaded into BI Accelerator.
- An index is created to this Info Cube and stored in BI Accelerator appliance. These are search engine indexes built using SAP's TREX search technology. They are stored in a file system using vertical decomposition. This results in highly compressed datasets that further contribute to fast processing speeds. When query is processed these indexes are loaded into memory. In memory, joins and aggregations are done at run time.
- At run time, query requests are sent to the analytic engine, which reroutes the query to the BI Accelerator and then to the end-user application

The BI Accelerator Server is a prepackaged hardware and software appliance that leverages SAP TREX search technology; combined with hardware technology developed by Intel Xeon 64-bit CPU blade servers and the blade servers must run on LINUX SLES 9.

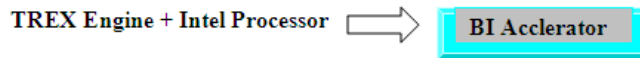


Figure2. TREX Engine

TREX engine is meant for search and classification of unstructured data, BI Accelerator is used to handle structured data.

- BI Accelerator engine and TREX are two different installations.
- BI Accelerator engine is part of the Analytic engine that manages the BIA index.
- The software allows reading, adding or changing data to the BIA index.
- The BI Accelerator optimizer is part of the BI accelerator engine that ensures the best possible read access to BIA index

A. WITH OUT BI ACCELERATOR

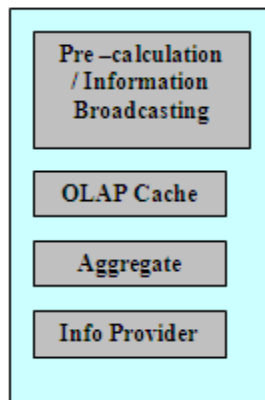


Figure3. With out BI Accelerator

SAP Net weaver BI checks the listed repositories to obtain the query result. Aggregates are Summarised (subset) physically stored views of data. OLAP Cache improves response time for only similar queries by caching query result sets and reading from cache instead of data base. Information Broadcasting: Run popular queries /Web/Work books in the background and push summary views of updated data to the users.

SAP Net Weaver BI Accelerator supplements traditional approaches (aggregates, query caching strategies, Information Broadcasting etc.) but does not replace them [3]. In many scenarios, SAP Net Weaver BI Accelerator may be superior to other approaches, but it is not mandatory to replace the current BI set-up and migrate to SAP Net Weaver BI Accelerator.

B. PERFORMANCE BOOST UP

BIA achieves impressive levels of data compression due to the column-based nature of its storage (blocks contain columns of data, not rows, sorted and tokenized making it easy to compress data of low cardinality), sorted and ordered and then stored in a way where the data is effectively the index. This is then horizontally partitioned across multiple blades (shared data storage, with fact data processed across all blade nodes but dimensions having a particular node affinity), and then loaded into RAM to create an in-memory database.

Apart from the reduced need for disk storage, column-based databases can work particularly well with data warehouses as a particular query, requesting just a small subset of the measures in a fact table, will request a much smaller amount of data from the cache as blocks contain just the data for a particular column, not the entire row with all the other measures that are not really required.

C. VERTICAL DECOMPOSITION

The table data is decomposed vertically by the accelerator engine, into columns that are stored separately. It makes more efficient use of memory space than row-based storage, since the engine needs to load only the data for relevant attributes or characteristics into memory. Hence only the relevant data columns are touched and accelerator engine can also sort the columns individually to specific entries to the top.

The column indexes are written to memory and cached as flat files. Thus input and output flows between in memory are smaller improving efficiency. This is appropriate for analytics, where most users want to see only a selection of data and attributes. In a conventional database, all the data in the table is loaded together, in complete rows.

Column Storage:

att1	att2	att3	att4	att5

Figure4. BI Accelerator stores tables by Column

To find instances of an attribute value:

- Go to the attribute column
- Read its row values

Classical DB Storage:

Tuple1			
Tuple2			
....			

Figure5. Vertical Decomposition of BIA Architecture

To find all instances of an attribute value:

- Go to first row
- Check the attribute value
- Go to next row
- Check the attribute value
- Repeat for each row in the table

D. HORIZONTAL DECOMPOSITION

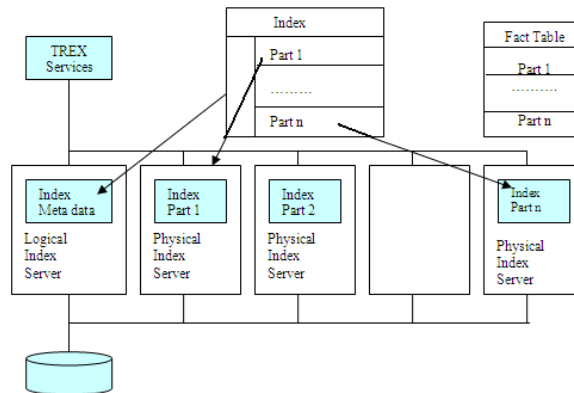


Figure6. Horizontal Decomposition of BIA Architecture

The BI Accelerator partitions the large fact tables as E and F 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.

E. BIACCELERATOR DEPLOYMENT

BI Accelerator fits seamlessly into existing environments. It can be activated for selected Info Cubes based on the requirement. Despite the BIA being an appliance companies can get considerable benefits only from optimized utilization.

- BIA accepts Info Cube data to be accelerated because BIA analytics will run more efficiently if they pass through the Dimension modeling of an Info cube.
- Identify the queries with high DB read times and check the scenarios with high data volume using the statistics (for example: ST03 or Admin cockpit) and include the corresponding info providers in BIA.
- If the landscape requires a high number of aggregates, the BI Accelerator reduces the overall TCO substantially.
- BW Accelerator implementations must be preceded with the below brief checks to identify : Impact and plans for Reporting DSO's, Impact and Plans for Optimizing BI Info Cubes for BWA, Impact of Info Cube modelling of BIA 'one-Time' and "annual maintenance' costs, Fast Track automated solutions for BIA Modelling.
- SAP Net Weaver BI Accelerator can be used in (almost) all cases where aggregates can be used.
- Remodeling of Info cubes can be undertaken with an intention of enhancing only the space and data load capabilities. Data load performance is also an additional benefit for considering re-modeling
- There is one exceptional scenario where SAP Net Weaver BI Accelerator and aggregates are not congruent: With Info Sets, an SQL statement is generated and, hence, the SAP Net Weaver BI Accelerator cannot be used.
- There are also some tools available that can assist for BIA implementations in automatic creation of Info Cube on a DSO, automatically modeling an Info Cube for BIA based on the requirement.
- BI support team should forecast the BIA data growth, physically view the indexed data.
- Support team should monitor space usage and alerts
- Need to forecast when next blades are required
- Review patch and upgrade.

F. TEST RESULTS

Performed several lab tests with real SAP Net Weaver BI customer data.

- Multi Providers with 6 Info cubes
- DSO's containing 800 million records together,
- Executed and tested 5 most important business critical queries.
- Developed a reconciliation report to reconcile the report data and the source data and did not find any data duplication or exceptions in the data.
- Performed data consistency check. Data was consistent.
- No changes have been made to the process chains.
- Realignment runs were faster.
- Index load was faster (56 121 954 records in 446 seconds).

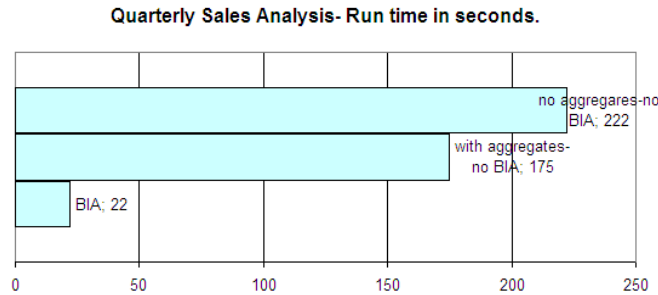


Figure7. Performance Comparison for Sales Report

III. CONCLUSIONS

BI Accelerator a very user friendly computer appliance which has preinstalled software on a predefined hardware and thus speeds up queries performance. Gives confidence to the customer that they can buy them selves out of performance bottlenecks by adding hard ware (blade racks). BI Accelerator is not a bullet that will resolve every query’s performance issue. We need to align expectations. BIA can provide exceptional query performance if applied in the right manner. Various points have been explained for BIA deployment and the tests reported in this paper indicate that adapting the accelerator offers practical benefits for SAP customers and IT Industries.

REFERENCES

- [1] Thomas Schroder,” *SAP BW Performance Optimization Guide*”, 1st ed., Galileo Press, Germany, (2006)
- [2] Andrew J. Ross,” *SAP Net weaver BI Accelerator*”, 1st ed., Galileo Press, Boston (MA), (2009)
- [3] Christian Merwald, Sabine Morlock,” *Data Ware Housing with SAP BW7, BI in SAP Net Weaver2004s*, 1st ed., Rocky Nook Inc., (2009)
- [4] Biao Fu, Henry Fu. *SAP BW a step by step Guide, 2002.*
- [5] Arshad Khan,” *SAP and BW Data Warehousing. How to plan and implement* (2005)
- [6] Danniell Knapp, ” *SAP Netweaver BI 7.0 Migration guide*“
- [7] http://www.sap.com/platform/netweaver_pdf/BWP_AR_IDC_BI_Accelerator.pdf
- [8] Kevin McDonald, Andreas Wilmsmeier, David C. Dixon, W.H Inmon, “*Mastering the SAP Business Information Ware house*” Wiley Publishing Inc, USA (2002)
- [9] Elizabeth Vitt, Michael Luckevich, Stacia Misner, “*Making better Business Intelligence Decisions faster*”, Microsoft Press, USA
- [10] Gary Nolan, Debasish Khaitan, “*Efficient SAP Netweaver BW implementation and Upgrade Guide.*