

Innovative Approach: SAP Netweaver Business Intelligence Accelerator



Applies to:

SAP Netweaver Business Intelligence Accelerator

Summary

This whitepaper is a brief summary about the SAP Netweaver Business Intelligence Accelerator. It gives an overview of the challenges faced prior to SAP BIA and how BIA helps in overcoming these challenges. It also tells about the installation requirements, initial results of BIA and some early BIA Customers.

Author: Denella D'souza

Company: Larsen and Toubro Infotech Limited

Created on: 20 January 2008

Author Bio



Denella D'souza works at Larsen and Toubro Infotech Limited as an SAP Business Intelligence Consultant.

Table of Contents

Introduction	3
Challenges	3
Primary Query Optimization Techniques	3
SAP NetWeaver Business Intelligence Accelerator	4
Architecture	4
Pre-Requisites	6
Using BIA Maintenance Index	7
Prerequisites	7
Process Flow	7
Scenario 1:	8
Scenario2:	8
Scenario3:	8
Estimating Memory Consumption by BIA Index	9
BI Accelerator Installment	10
Initial BI Accelerator Lab Test Results (2005)	11
Some Early BI Accelerator Customers	12
Related Content	12
Disclaimer and Liability Notice	13

Introduction

The demands of the business analytics market have been increasing and this has driven the requirement for quick decision making based on the current data. The two major demands are Speed and Flexibility wherein the data should be available at the right time and at a faster rate. Over the years, several enterprises have implemented data warehouses for business intelligence with the fact of gaining some profits over their investments, for customer satisfaction and quick decision making processes. Still all their demands still have remained unsatisfied. At the same time they also have to keep up with increasing data volume. To satisfy such demands at a go, SAP has brought forth SAP BI Accelerator.

The BI Accelerator appliance benefits businesses handling high volume of data and meets all the new and growing enterprise requirements. This whitepaper gives an overview of the need of BI Accelerator and its benefits.

Challenges

- **Increasing Data Volume:** The data volume continues to grow as new data comes in over time, for example, the data for RFID, EDW etc. Even with the increase in data volume the information should be available at the speed of thought.
- **Speed and Flexibility:** The information should be available at the right time and at a faster rate.
- **Increasing Number of Users:** More business users and casual users work on the system. Therefore, there should be quick and easy scalability.
- **Additional Administration Efforts:** These include various administration challenges. Potential performance issues have to be taken into consideration.

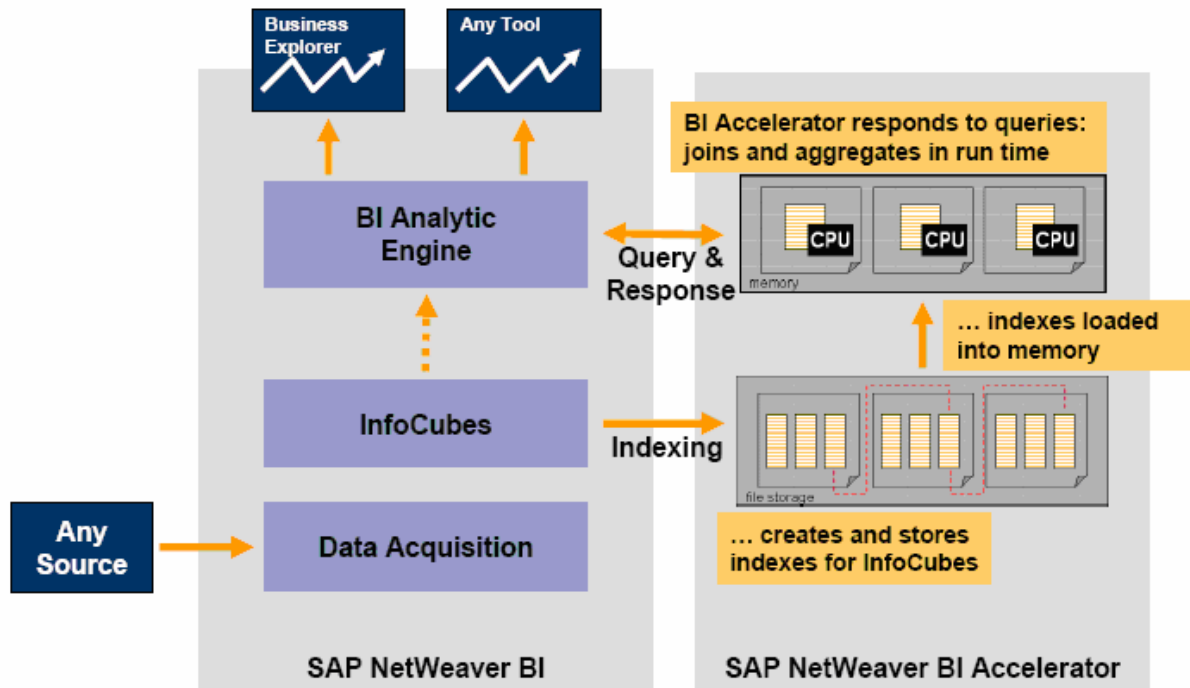
Primary Query Optimization Techniques

- **Aggregates:** Relational Aggregates are wherein we have a subset of the Infocube data and then use these subsets into queries to speed up query performance. Therefore, these are BI objects designed solely to speed up the query performance. They structurally look same as an Infocube. Basically, it helps to reduce the volume of data read. To improve the query performance, you need to decide or pre-calculate aggregates for main queries and results. But using relational aggregates has a trade-off between performance and flexibility. Aggregates cannot be created for all possible queries. Thus, the typical strategy is to create general aggregates that more queries can utilize, but the performance benefit per query is less significant.
- **OLAP cache:** It accelerates response for similar queries by caching query result sets., and reading from cache instead of database.
- **Reporting agent/information broadcasting:** Run popular queries in the background and push summary views of updated data to users.

SAP NetWeaver Business Intelligence Accelerator

SAP NetWeaver Business Intelligence accelerator represents an innovative new approach for boosting query performance. It is based on TREX (Text retrieval and Extraction) which is SAP's own developed search engine technology. Due to this the performance is typically 10 times faster or even 100 times or more. All this is possible without the changing the front-end tools that users have utilized when they work with the BI System.

Architecture





- The BI Accelerator is an applicant type model. It connects to the SAP Netweaver Business Intelligence system serving the sole purpose of speeding up queries. As an appliance model it connects to the SAP Netweaver Business Intelligence System. We need to decide for which infocube we need to create indexes which will be stored in the BI accelerator Appliance.
- When a job is run, the data is copied into the search engine indexes and is made available in the BI Accelerator as data in these indexes is activated into the memory. When queries run against SAP Netweaver BI System, any tool that you use with SAP Netweaver BI System is viable. The analytical engine knows if a cube has been turned on for BI Accelerator for which the query request should shift to the BI Accelerator for processing.
- The read operation for the query takes place into the memory on the BI Accelerator Appliance and the performance is very fast and it is highly parallelized.
- The data is loaded into the infocube in the same way as in the SAP Netweaver BI system. In the BI accelerator, the data needs to be copied from the infocube to the BI accelerator system using roll-up process.

In short, BI accelerator essentially replaces relational aggregates.

- BI Accelerator can benefit businesses that have high volumes of data. The read performance of BI queries is greatly improved using BI Accelerator. This new tool enables 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 volume and high frequency of queries.
- BI Accelerator is also useful when aggregates or database indexes are not sufficient, or when these methods become too complex or too costly to maintain. For example, it may not be practical or cost effective to create and maintain all the aggregates that are needed to support performance for all queries for an InfoCube.
- There are two major components in the BI Accelerator architecture that contribute to the high performance.
 1. BI Accelerator Engine
 2. BI Accelerator Indexes
- There is increased end-user satisfaction with consistently high query performance.
- Significant reduction in the total cost of ownership: low maintenance, high scalability.
- The performance gain using BI Accelerator can be very great, boosting the response time by a factor of 10-100

Pre-Requisites

IBM SYSTEMS SOLUTION FOR BI ACCELERATOR FROM SAP

BI accelerator software provided by SAP based on SAP NetWeaver 2004s

- running under 64-bit Novell® SUSE® Linux®

IBM BladeCenter Chassis (with redundant HW layout)

- Ethernet Switch (customer configurable)
- Fiber Channel Switch (fix)

IBM BladeCenter HS20 Blade

- 2 Intel Xeon® EM64T CPU
- 8 GB RAM
- Fiber Channel Adapter

IBM System Storage DS4300

- I/O through FiberChannel (separate I/O from network)
- Raw Disk Sizes between 1 TB to 3.5 TB
- RAID setup for performance and High Availability

IBM General Parallel File System (GPFS)

- Performance optimized shared filesystem for Linux Cluster
- HA configuration for minimized downtime

- SAP NetWeaver 7.0 TREX is the underlying engine for the BI accelerator. The data of the BI InfoCubes are replicated to the TREX engine and are stored as TREX indices. The BI accelerator clearly reduces the response time, especially for large data volumes.
- BI accelerator can only in connection with the *SAP NetWeaver 7.0 BI* starting with *SAP NetWeaver 7.0 Support Package Stack (SPS) 5*. BI accelerator is a special Linux 64 bit version of *SAP NetWeaver 7.0 Search and Classification (TREX)* which is delivered on preconfigured hardware. BI accelerator cannot be used for the regular TREX search and classification functionality.
- The BI Accelerator is installed on a preconfigured blade system. A blade system consists of hosts in the form of server blades. The server blades are connected to central disk storage. This is referred to here as a file server, regardless of the underlying hardware.
- The special feature of a BI Accelerator installation on a blade system is that the BI Accelerator software can be stored centrally as well as the BI Accelerator data. This means that the software will be installed only once on the file server. Maintaining the system is efficient because you only have to implement software updates once. All server blades on which BI Accelerator is running access the same program files. However, each server blade has its own configuration files. The configuration files in the directory <TREX_DIR> are only used as templates. A script creates a separate subdirectory for each server blade and copies the configuration files to this subdirectory.

Using BIA Maintenance Index

For each Infocube, Only one SAP Netweaver BI Accelerator Index can be made. The BI accelerator Index contains all the data from the Infocube. The Application Logs function key in the BIA index maintenance wizard helps to choose the particular processes for which the log needs to be displayed.

The different processes available are as follows:

- Initial Filling
- Roll up
- Compress Infocube
- Delete request
- Change Run
- Check

If a BI Accelerator index is available, the Maintain BIA index Properties dialog box appears when the BIA Index Properties function key is clicked. This helps us to make certain settings like it provides us an option of storing BIA index data completely in the main memory. This setting is advisable only if there is enough memory available and the index is used frequently. Another option provided is to set or change the status of the BIA index (active or inactive).

Prerequisites

Communication between the BI system and the BI accelerator server takes place using the RFC modules. Certain settings are required to connect the BI accelerator server to a BI system. To set the RFC connection, the RFC destination needs to be specified in the transaction RSADMIN. The list of RFC Destinations can be seen in the transaction SM59. The RFC Destination specified in RSADMIN has to be present in the list of RFC destinations in SM59.

Process Flow

Access from Data warehousing Workbench:

In the Data warehousing Workbench, under the Modeling functional area in the navigation window, we can choose the Infoprovider which contains the Infocube on which we want to create index. In the context menu of the Infocube, choose Maintain BI Accelerator Index. The BIA index maintenance wizard appears.

Access from transaction RSDDV:

We can also use transaction RSDDV to access the BIA index maintenance wizard. On the Aggregate/BI accelerator index, select Infocube screen and select the required infocube. Then select the BIA index.

There are two scenarios in which the BIA index can be created.

Scenario 1:

The BIA index wizard is called for an infocube that does not yet have a BI accelerator index.

Step 1: Creating a BI accelerator index

When this step is executed, a BIA index is created for that particular infocube as long as it does not have indexes created by another BIA. These tables contain fact and dimensions of the Infocube as well as master data tables that contain SIDs, the S, X and Y tables of the info objects. A logical index is also created. This contains the metadata of the BI accelerator. Finally, the system activates the BI accelerator index.

Now to use the BI accelerator index in reporting, we will have to fill it up with data first.

Step 2: Filling a BI accelerator index

Then we have to schedule a job to run in background processing (RSDDTREX_AGGREGATES_FILL) and choose Save. When this step is executed the system reads the data in the tables of the Infocube star schema from the database and writes them into the corresponding indexes on the BI accelerator server. If the master tables already have indexes created and filled by another BI accelerator index, then only those records that have been subsequently added have to be indexed (read mode/ fill mode 'D' during indexing)

Step 3: Completing BI accelerator index maintenance

After the BI accelerator has been filled, you can choose Continue to continue to the first part of BI accelerator index maintenance. The BI accelerator index is available and can be used for queries.

Scenario 2:

BIA Index wizard is called for an infocube which has a BIA index that is already filled with data.

The BIA index already created can be either deleted or can be deactivated temporarily. If the index is deleted then the system deletes the definition and the settings of the BI accelerator index in the BI system and the logical index and all indexes for the tables of the enhanced star schema of the Infocube on the BI accelerator server.

Scenario 3:

BIA index wizard for an infocube can be called that has already active BI accelerator index, but has not yet been filled or completely filled with data.

Since an active BI accelerator index that can be used for reporting is already available, you can either continue to fill it with data or delete it at this point. To continue filling data we have to select the option 'Continue Filling'. To delete the BI accelerator index, select 'Delete'.

Estimating Memory Consumption by BIA Index


SAP Note 917803

- Implement and run the program that is attached to the note i.e.
 - [zz_biamemconsumption_bw3x_sub.txt](#)
 - [zz_biamemconsumption_bw3x.txt](#)
- Possible long runtime, so schedule over weekend
- You can find the output of the program in the spool list in the form as shown below:


TableName	Memory [kB]	Attribute
/B10/D0BWVC_0032	4	DIMID SID_OVC_PROD2
/B10/D0BWVC_0033	37	DIMID SID_OVC_SELL2
/B10/D0BWVC_0034	11	DIMID SID_OVC_LOC2
/B10/D0BWVC_0035	5	DIMID SID_OVC_REDU2
/B10/D0BWVC_003P	3	DIMID SID_ORECORDTP
/B10/D0BWVC_003T	16	DIMID SID_OFISCVARNT
/B10/D0BWVC_003U	3	DIMID SID_OCURRENCY
/B10/SCURRENCY	6	CURRENCY SID
/B10/SFISCPER	170	FISCVARNT FISCPER SID
/B10/SFISCVARNT	2	FISCVARNT SID
/B10/SFISCYEAR	8	FISCVARNT FISCYEAR SID
/B10/SUNIT	7	UNIT SID
/B10/SVC_COL2	3	VC_COL2 SID
/B10/SVC_GROU2	2	VC_GROU2 SID
/B10/SVC_HOME2	9	VC_HOME2 SID
/B10/SVC_LOC2	10	VC_MARK2 VC_LOC2 SID
/B10/SVC_MARK2	3	VC_MARK2 SID
/B10/SVC_MAT2	3	VC_MAT2 SID
/B10/SVC_PAYM2	3	VC_PAYM2 SID
/B10/SVC_REDU2	2	VC_REDU2 SID
/B10/V0BWVC_003F	98	KEY_OBWVC_003P
/B10/XVC_PERS2	22	SID S__OVC_HOME2
/B10/XVC_PROD2	67	SID S__OVC_GROU2
/B10/XVC_UNIT2	5	SID S__OVC_HOME2
/B10/YVC_PROD2	70	SID DATETO DATEFROM
* 569		

BI Accelerator Installment


■ **The BI Accelerator Appliance**



BIA appliance by IBM




POWERED BY
SAP NetWeaver™





Storage Subsystem
SUSE Linux Enterprise Server

**Blade
Chassis / Enclosure**



CPUs by Intel

BIA appliance by HP

BI Accelerator is only available as an appliance (“in a box”).

Blade servers with 64 bit Intel Xeon CPUs.

Operating System (as of today):

- (a) SUSE Linux SLES, 64 bit
- (b) No support for other OS planned.

Optimized storage included.

Hardware partners (as of today):

- (a) Hewlett Packard
- (b) IBM

(No restrictions for the platform for the SAP NetWeaver BI implementation)


Initial BI Accelerator Lab Test Results (2005)

This was done with real SAP NetWeaver BI customer data i.e.


- Multiprovider with 9 infocubes.
- About 850 million records altogether.
- Customer Aggregates used.
- The 7 most important and critical queries tested.
- No blade appliance in place, basic server utilized.

Initial Test results

Global Sales – Comparative Runtime in Seconds



Product Margin Reporting – Comparative Runtime in Seconds



© SAP AG 2006, <#>

THE BEST-RUN BUSINESSES RUN SAP™



=>Average Improvement factor: 25x faster.

Some Early BI Accelerator Customers



Related Content

[SAP Library - Performance Optimization with SAP NetWeaver BI Accelerator](#)

[Know How Network: The SAP NetWeaver Business Intelligence Accelerator by Ron Silberstein](#)

[SAP Network Blog: What is BI Accelerator? WebLog 16 Mar 2006 Irvati Andhare](#)

Disclaimer and Liability Notice

This document may discuss sample coding or other information that does not include SAP official interfaces and therefore is not supported by SAP. Changes made based on this information are not supported and can be overwritten during an upgrade.

SAP will not be held liable for any damages caused by using or misusing the information, code or methods suggested in this document, and anyone using these methods does so at his/her own risk.

SAP offers no guarantees and assumes no responsibility or liability of any type with respect to the content of this technical article or code sample, including any liability resulting from incompatibility between the content within this document and the materials and services offered by SAP. You agree that you will not hold, or seek to hold, SAP responsible or liable with respect to the content of this document.