

How to Analyze Query Performance through BI Accelerator



Applies to:

BI NetWeaver 2004s For more information, visit the [Business Intelligence homepage](#).

Summary

This paper describes How to analyze Query performance through BI Accelerator

Author: Deshmukh Bharti

Company: Accenture Services Private Ltd

Created on: 19th of April 2010

Author Bio

Deshmukh Bharti is working as SAP BI Consultant in Accenture Services Private Ltd and having extensive experience in implementation of BI projects.

Table of Contents

1.) Introduction.....	3
2.) Pre-Requisites	3
3.) SAP BI Accelerator Architecture.....	4
4.) SAP BI Accelerator at Work	5
5.) SAP BI Accelerator Query Processing Steps.....	8
6.) Query execution time without BIA indexing.....	9
Step 1.....	9
Step 2.....	9
Step 3.....	10
Step 4.....	10
7.) Query execution time with BIA indexing.....	11
8.) Benefits and Shortcomings of SAP BI Accelerator	12
Benefits	12
Shortcomings	12
Related Contents	13
Disclaimer and Liability Notice.....	14

1.) Introduction

The BI accelerator is a highly scalable analytic server that processes queries initiated by users of SAP Net Weaver BI. Its uniqueness and the features that make BI accelerator highly scalable is the use of SAP's TREX search technology in conjunction with blade server architecture provided by its hardware partners.

SAP BI Accelerator comes into picture to address serious performance problems of BI queries. BI Accelerator is a plug-in appliance delivering:

- Scalability in terms of expandable blade hardware infrastructure delivered by IBM or HP
- Performance in memory processing, fully leveraging the power of Intel Processors
- Flexibility in terms of no dedicated tuning efforts and reduced load times.

2.) Pre-Requisites

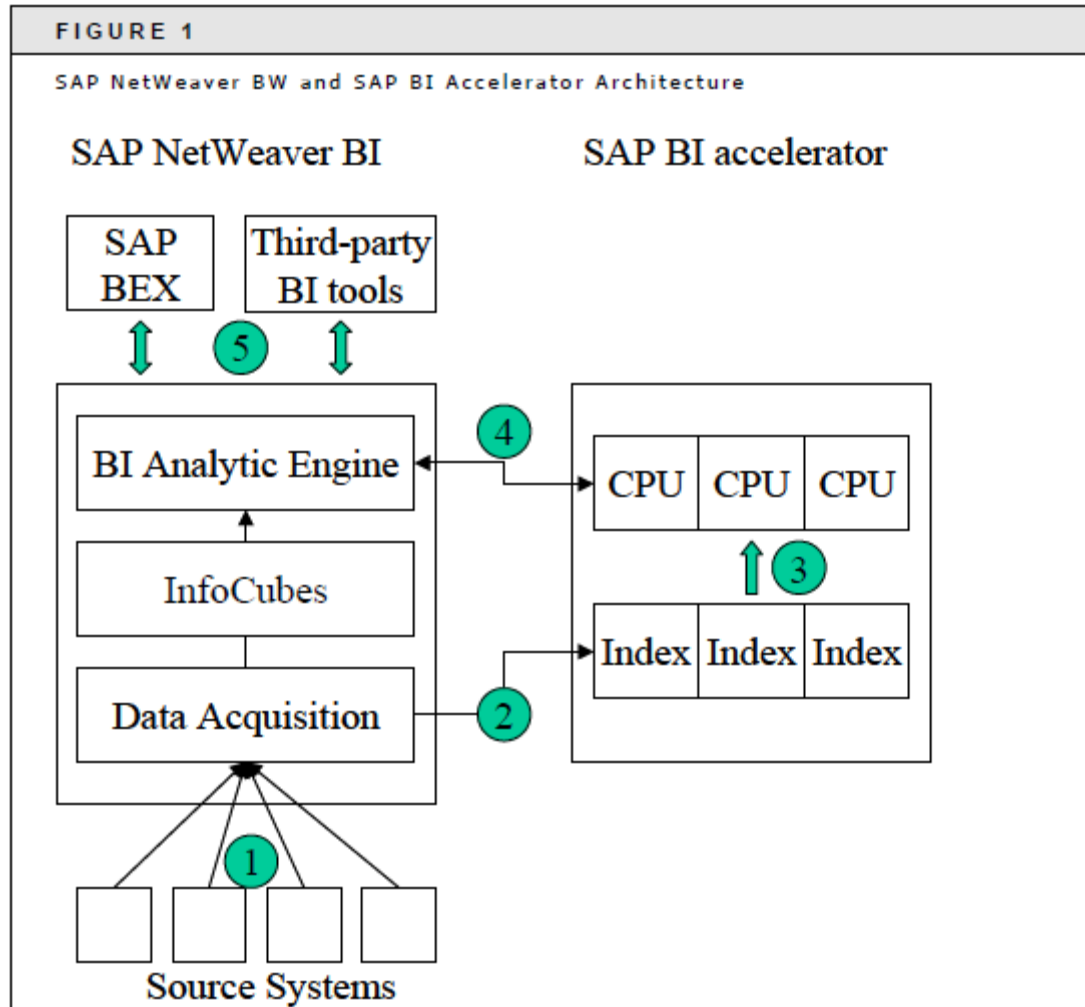
The requirements for BI Accelerator are:

- SAP BI in SAP NetWeaver 2004s.
- We also require an installation based on 64 bit architecture as BI Accelerator is based on TREX technology.
- Queries with high database read times.
- Ad-hoc reporting on high data volumes.

Note: BI Accelerator cannot be used for standard KMC functionality. A TREX installation cannot be used as a BI Accelerator

3.) SAP BI Accelerator Architecture

A high-level BI Accelerator architectural diagram is shown on the right side of Figure 1.



SAP BI accelerator includes indexes that are vertically inverted reproductions of all the data included in InfoCubes (i.e., fact and dimension tables as well as master data). Note that there is no relational or other database management system in BI accelerator. There is only a file system, and indexes are essentially held as flat files.

The second primary component of SAP BI accelerator is the engine that processes the queries in memory. The software is running on an expandable rack of blade servers. The operating system used for BI Accelerator is 64-bit Linux, so in addition to having no database license cost, there is also no OS license cost.

4.) SAP BI Accelerator at Work

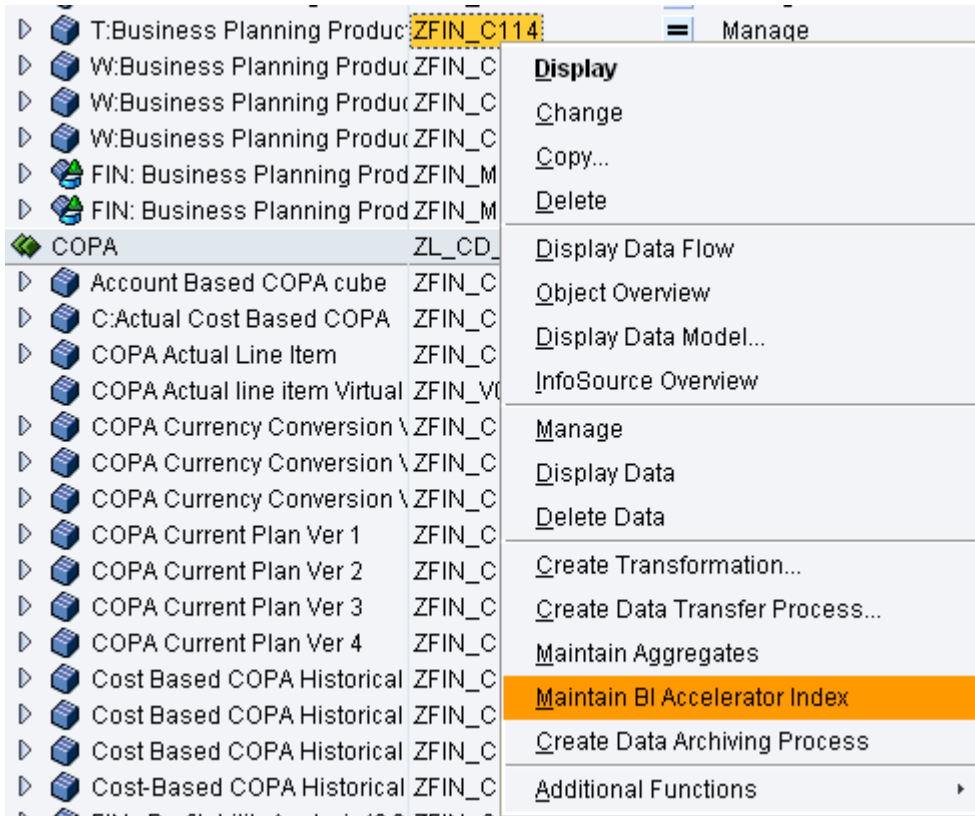
1. Data is loaded from source systems into an SAP InfoCube.

The screenshot displays the SAP BI Accelerator interface. On the left, a tree view shows various data targets, with 'T:Actual Cost Based COPA' (ZFIN_C055) selected. The main window is titled 'Selectable Data Targets for Administration' and shows a table with columns: Name, D..., Technical Name, and Table Type. The selected target is 'T:Actual Cost Based COPA' with Technical Name 'ZFIN_C055' and Table Type 'InfoCube'. Below this, there are tabs for 'Contents', 'Performance', 'Requests', 'Rollup', 'Collapse', and 'Reconstruction'. The 'Requests' tab is active, showing a table of requests for the selected InfoCube. The table has columns: Requ..., R..., C..., C..., R..., Re..., Loa..., DTP/InfoPackage, Request D..., Update Date, and Selection. The requests are listed with their IDs, status icons, and details.

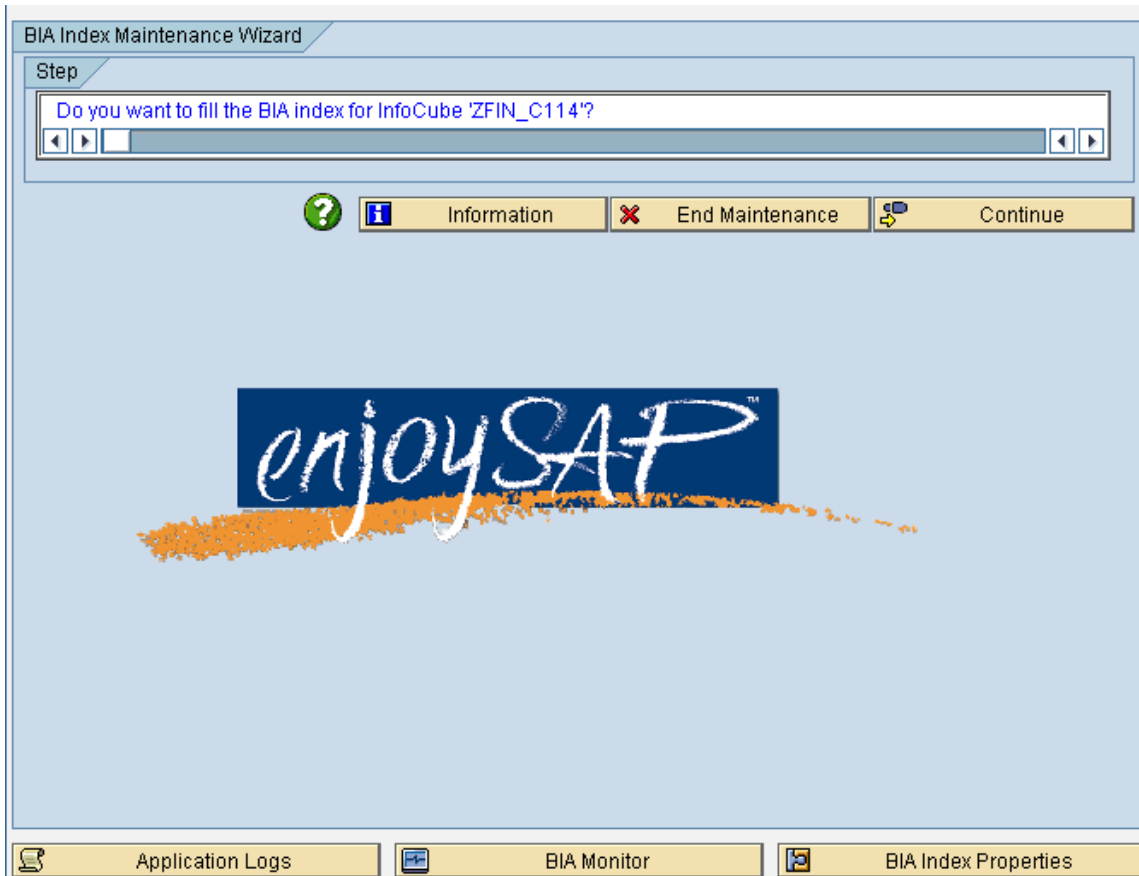
Requ...	R...	C...	C...	R...	Re...	Loa...	DTP/InfoPackage	Request D...	Update Date	Selection
24599...		✓	✓	○	○	○	T: DELTA/8ZFING055/ZFI...	19.03.2010	19.03.2010	
24592...		✓	✓	○	○	○	T: DELTA/8ZFING055/ZFI...	19.03.2010	19.03.2010	
24575...		✓	✓	○	○	○	T: DELTA/8ZFING055/ZFI...	18.03.2010	18.03.2010	
24568...		✓	✓	○	○	○	T: DELTA/8ZFING055/ZFI...	18.03.2010	18.03.2010	
24549...		✓	✓	○	○	○	T: DELTA/8ZFING055/ZFI...	17.03.2010	17.03.2010	
24542...		✓	✓	○	○	○	T: DELTA/8ZFING055/ZFI...	17.03.2010	17.03.2010	
24524...		✓	✓	○	○	○	T: DELTA/8ZFING055/ZFI...	16.03.2010	16.03.2010	
24516...		✓	✓	○	○	○	T: DELTA/8ZFING055/ZFI...	16.03.2010	16.03.2010	
24498...		✓	✓	○	○	○	T: DELTA/8ZFING055/ZFI...	15.03.2010	15.03.2010	
24491...		✓	✓	○	○	○	T: DELTA/8ZFING055/ZFI...	15.03.2010	15.03.2010	
24474...		✓	✓	○	○	○	T: DELTA/8ZFING055/ZFI...	14.03.2010	14.03.2010	
24467...		✓	✓	○	○	○	T: DELTA/8ZFING055/ZFI...	14.03.2010	14.03.2010	
24450...		✓	✓	○	○	○	T: DELTA/8ZFING055/ZFI...	13.03.2010	13.03.2010	
24443...		✓	✓	○	○	○	T: DELTA/8ZFING055/ZFI...	13.03.2010	13.03.2010	

2. An index is built for this InfoCube and stored inside the BI accelerator appliance. These are search engine indexes built using SAP's TREX search technology. They are stored in a file system (not a database system) using vertical decomposition (a column-based approach as opposed to the row-based approach that requires more read time). This results in highly compressed data sets that further contribute to fast processing speeds.


To create index on a cube, from context menu of that cube select tab strip "Maintain BI Accelerator index"



A BIA Index Maintenance Wizard will be appeared then select Continue button

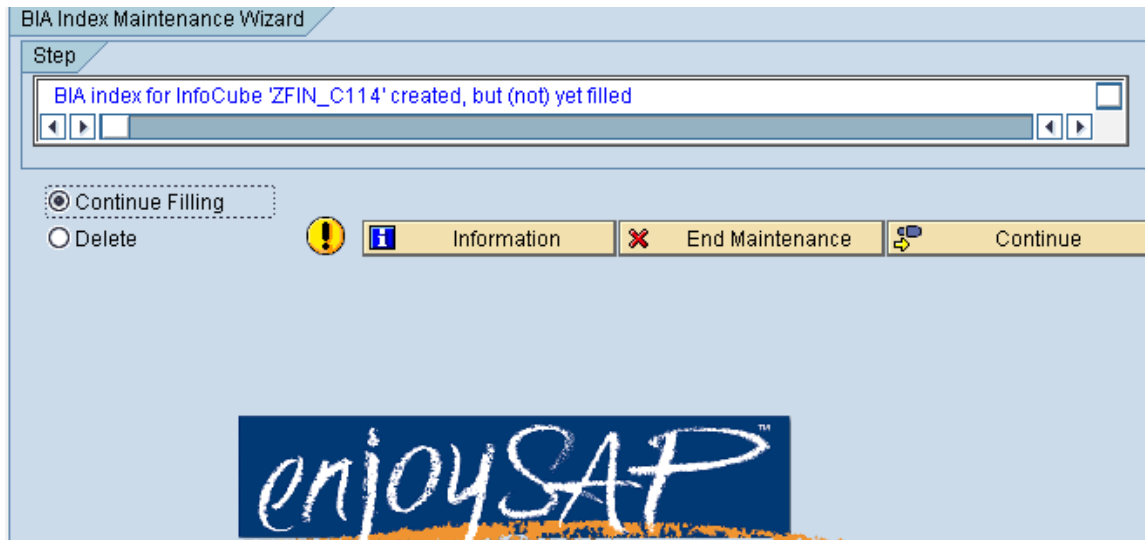


Following message will be appear to continue again

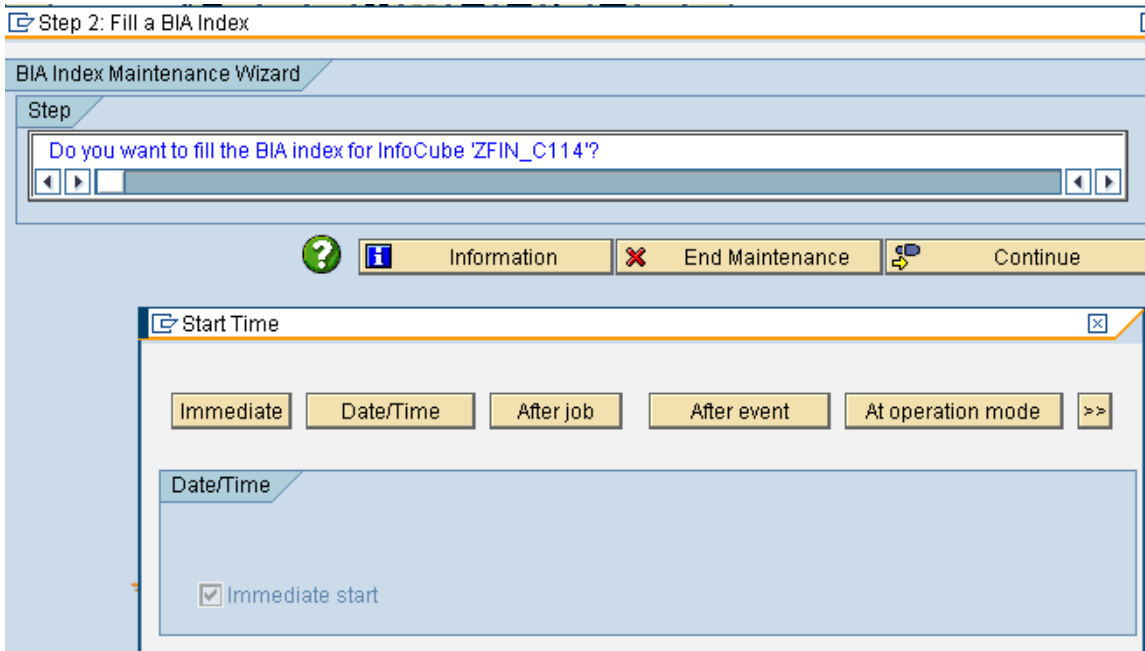


BIA index for InfoCube 'ZFIN_C114' created, but (not) yet filled

Press again continue button to fill the index.




Now to fill BIA index we can start the job immediately or as per business requirement.



Then save it



A job RSDDTREX_AGGREGATE_FILL will be scheduled.



We can monitor in SM37 until the job finished.

Job	Spool	Job Doc	Job Created	Status	Start date	Start time	Duration(sec.)	Delay (sec.)
BW_TR_RSDDTREX_AGGREGATES_FILL			LL5774	Finished	15.04.2010	13:40:11	28	1
*Summary							28	1

Now the Index has been created on cube.

3. BI accelerator indexes are loaded into memory where the query is processed. In memory, joins and aggregations are done at run time. Loading of indexes into memory happens automatically at first query request, or it can be set for preloading whenever new data is loaded.
4. At run time, query requests are sent to the analytic engine, which reroutes the query to the BI accelerator.
5. Query results are returned to the end-user application.

Note: Steps 1 through 3 above typically are performed offline, (e.g., during less critical times). Steps 4 and 5 below are executed at actual query time.

5.) SAP BI Accelerator Query Processing Steps

To describe how queries are processed by SAP BI accelerator, it is first worthwhile to describe how they are processed within the traditional SAP NetWeaver BI architecture. The left side of Figure 1 depicts this SAP NetWeaver BI architecture. In this case, the steps are:

1. Query is launched from SAP Business Explorer (BEX) or a third-party BI tool.
2. Query evaluates whether there is a pre calculated data set (usually calculated during off-hours). If one exists, the query retrieves data from that data set.
3. If a pre calculated template does not exist, the query checks the OLAP Cache for the necessary data. The OLAP Cache doesn't benefit the first person launching the query, but will benefit all subsequent requests for that same query.
4. If the required data does not exist in the OLAP Cache, then the query looks for aggregate tables or materialized views that may exist. These pre aggregated views are not as fast as processing the query against pre calculated data sets or OLAP Cache, but they are still faster than going against the final layer, the InfoProvider.
5. The final option to execute the query is to run it against the InfoProvider, in this case the SAP InfoCube. These results in the slowest processing times as compared to the other three choices listed above.

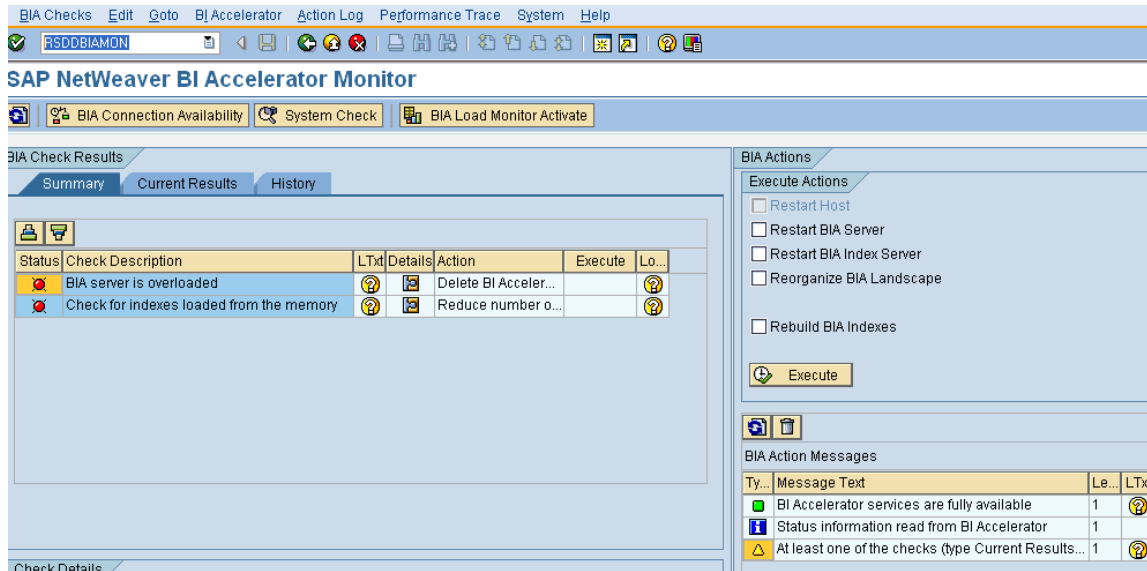
Executing the same query with SAP BI accelerator results in a somewhat different set of steps:

Here analysis has done through query execution time with and without BIA indexing.

6.) Query execution time without BIA indexing

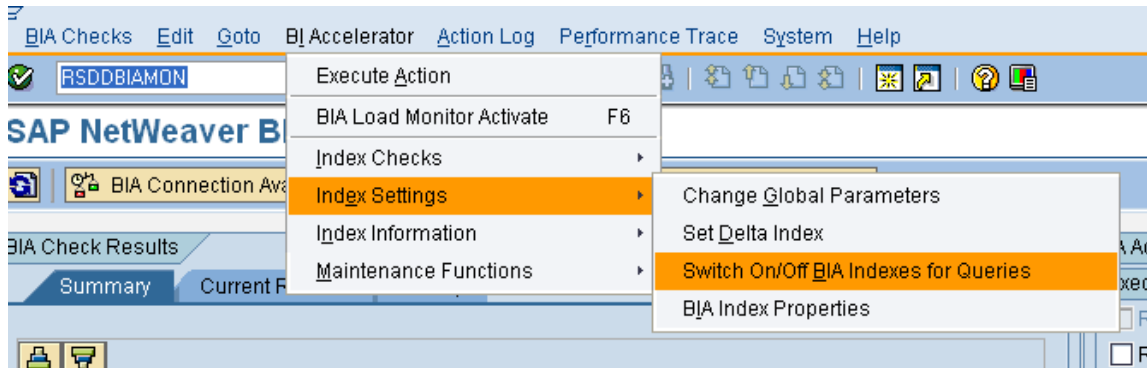
Step 1.

In tcode RSDDBIAMON, switch off the indexing of cube




Step 2.

Go BI Accelerator button from menu and switch off index for query.



Check right mark against that cube on which query is needs to be executed



and save it. 

Step 3.

Through tcode RRMX execute the query based on cube as given in above example.

The screenshot shows the RRMX transaction screen with the following filter criteria:

- Material: = [] To []
- PO Plan Month: = [] 001.2007 To [] 012.2007
- Mfg Order Plant: = [] To []

At the bottom, there is a toolbar with buttons for Execute, Check, Save, Copy, Paste, Delete, Print, and Cancel (F12).

Note: make sure time interval should be good range to fetch the data from table of cube rather than cache memory.

Step 4.

Now total execution time can be capture through workload monitor tcode ST03

Drilldown BI workload tab strip

Workload Monitor

The screenshot shows the Workload Monitor (ST03) interface. The left sidebar shows a tree view with 'BI Workload' selected. The main area displays a table titled 'Overview of instances and last analysis data'.

ABAP Instance Name	Host Name	Last Measurement	Time	Active
640_TOTAL	640_TOTAL	31.03.2010	23:59:56	
640_bsdudx05_CRB_00	640_bsdudx05_CRB_00	01.04.2010	18:59:43	
640_bsdudx20_CRB_10	640_bsdudx20_CRB_10	01.04.2010	18:59:44	
640_bsdudx20_CRB_20	640_bsdudx20_CRB_20	01.04.2010	18:59:34	

Then double click on Today list button and then in right hand side select Query from Aggregation.

Workload in System CRB

Query	No. of Nav.	Total time	% of Total
TOTAL	23	48	2.947,7
ZFIN_M050	2	5	392,5
ZWVM_M065	1	2	323,7
ZWVM_M090	1	2	254,6
ZBP_M300	2	7	229,2

Then check the total time of executed query from below list

InfoCube	Name of Query	No. of Nav.	Total time	% OLAP	% DB	% Frontend	% Plan	% Un	OLAP Time	DB Time	Frontend	Planning	Unass. Time
TOTAL	TOTAL	48	2.947,7	73,98	24,22	1,80	0,00	0,00	2.180,7	714,0	53,0	0,0	0,0
ZWVM_M065	ZCD_WM_GLBL_ZWVM_M065_Q0003	2	323,7	72,81	25,64	1,58	0,00	0,00	235,7	83,0	5,1	0,0	0,0
ZWVM_M090	ZCD_WM_GLBL_ZWVM_M090_Q0001	2	254,6	71,88	27,42	0,71	0,00	0,00	183,0	69,8	1,8	0,0	0,0
ZFIN_M050	ZCD_FIN_GLBL_ZFIN_M050_Q0045	1	247,5	99,15	0,85	0,20	0,00	0,00	245,4	1,6	0,5	0,0	0,0
ZBP_M300	YCD_BP_GLBL_ZBP_M300_Q0002_TST	6	222,5	44,85	54,38	0,81	0,00	0,00	99,8	121,0	1,8	0,0	0,0
ZFIN_M130	ZCD_FIN_GLBL_ZFIN_M130_Q0001	5	200,8	95,47	3,04	1,49	0,00	0,00	191,7	6,1	3,0	0,0	0,0
ZRTP_M001	ZCD_RTP_GLBL_ZRTP_M001_Q0002	1	199,2	34,34	65,51	0,15	0,00	0,00	68,4	130,5	0,3	0,0	0,0
ZFIN_M020	ZCD_FIN_GLBL_ZFIN_M020_Q0001	3	196,0	92,35	7,55	0,10	0,00	0,00	181,0	14,8	0,2	0,0	0,0
ZSTC_M100	ZCD_STC_GLBL_ZSTC_M100_Q0001	1	153,1	18,29	81,19	0,52	0,00	0,00	28,0	124,3	0,8	0,0	0,0
ZFIN_M050	ZCD_FIN_BOBJ_ZFIN_M050_Q0001	4	145,0	80,48	3,72	15,79	0,00	0,00	116,7	5,4	22,9	0,0	0,0
ZMFG_M020	ZCD_MFG_GLBL_ZMFG_M020_Q0002	1	136,3	97,36	2,49	0,15	0,00	0,00	132,7	3,4	0,2	0,0	0,0

Here No. of Navigation is showing how many time that query have been executed.

Double click on that query name and get all statistics.

7.) Query execution time with BIA indexing

Perform same step from 1 to 4 as in without indexing but here now indexing will be switch on in step 1.

InfoCube	Switched Off for Q...	Last Chan...	Time St...
ZFIN_C103	<input type="checkbox"/>	LL5774	20.100...
ZFIN_C104	<input type="checkbox"/>	LL5774	20.100...
ZFIN_C105	<input type="checkbox"/>	LL5774	20.100...
ZFIN_C114	<input type="checkbox"/>	LL5774	20.100...
ZFIN_C115	<input type="checkbox"/>	LL5774	20.100...
ZFIN_C200	<input type="checkbox"/>	LL5774	20.100...
ZFIN_C201	<input type="checkbox"/>	EP3774	20.100...

Then we can capture query execution time with and without indexing.

Individual Records of Query ZCD_FIN_GLBL_ZFIN_M100_Q0001

InfoCube	Name of Query	Frontend	Total time	OLAP Time	% OLAP	DB Time	% DB	Frontend	% Frontend	Planning	% Plan	Unass.Time	% Un
ZFIN_M100	ZCD_FIN_GLBL_ZFIN_M100_Q0001	BEX3	35,7	23,9	66,95	1,5	4,20	10,3	28,85	0,0	0,00	0,0	0,00
ZFIN_M100	ZCD_FIN_GLBL_ZFIN_M100_Q0001	BEX3	5,2	3,3	63,46	1,0	19,23	0,9	17,31	0,0	0,00	0,0	0,00

Here some statistics showing how the BI Accelerator improved query performance.

InfoCube name	Time to create BIA Index sec	Number of Records in F table	Query name	Time without BWA index	Time for Same Query with BWA Index	Query Performance improvement in %
ZFIN_C051	6533	30462712	ZCD_FIN_GLBL_ZFIN_M050_Q0026	45.1	18.4	145.11%
ZFIN_C053	5583	81906293	ZCD_FIN_GLBL_ZFIN_M050_Q0026	66.9	27.9	139.78%
ZFIN_C054	7766	107767126	ZCD_FIN_GLBL_ZFIN_M050_Q0026	21.7	19	14.21%
ZFIN_C055	107767126	7702484	ZCD_FIN_GLBL_ZFIN_M050_Q0026	14.4	7.3	97.26%
ZFIN_C056	1486	10944364	ZCD_FIN_GLBL_ZFIN_M050_Q0026	14.6	9.1	60.44%
ZFIN_C101	510	9092794	ZF5_FIN_GLBL_ZFIN_M100_Q0004	36	8.6	318.60%
ZFIN_C102	431	2850938	ZF5_FIN_GLBL_ZFIN_M100_Q0004	17.7	6.1	190.16%
ZFIN_C103	1,774	23490066	ZF5_FIN_GLBL_ZFIN_M100_Q0004	46.6	4.8	870.83%
ZFIN_C104	63	0	ZF5_FIN_GLBL_ZFIN_M100_Q0004	4.5	3.3	36.36%
ZFIN_C105	203	2193417	ZF5_FIN_GLBL_ZFIN_M100_Q0004	10	2.6	284.62%
ZFIN_C115	93	139669	ZF5_FIN_GLBL_ZFIN_M100_Q0004	2.6	2	30.00%

8.) Benefits and Shortcomings of SAP BI Accelerator

Benefits

- Faster query processing and response time
- Faster load times, as aggregate change runs due to master data changes are handled by the BI accelerator rather than on top of InfoCubes
- Lower maintenance costs:
 - BI accelerator eliminates the need to create relational aggregates.
 - BI accelerator may eliminate the need to deal with an OLAP Cache.
 - BI accelerator may decrease the need for logical partitioning on the NetWeaver BI side. However, there are other benefits beyond improving query processing speeds to having logical partitions.
 - BI Accelerator results in less planning and tuning on the part of DBAs. Attractive packaging as an appliance that is preconfigured for analytic processing using SAP software and partners' hardware, which allows nonintrusive implementation
- High potential scalability, as demands grow, system scales up by adding blades

Shortcomings

- Currently the data source for BI accelerator can only be an SAP InfoCube. It does not work with other SAP data sources such as ODS.
- There's currently a one-to-one relationship between an instance of SAP NetWeaver BI and a BI accelerator. Sharing of multiple SAP NetWeaver BI instances with a single BI accelerator is not yet supported.

Related Contents

For more information, visit the [Business Intelligence homepage](#).

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.