IBM Servers and Storage Demos
IBM SAP Business Intelligence Accelerator

Abstract
SAP's Business Intelligence Accelerator (BIA) has been providing customers with super fast query processing to meet their business intelligence needs. Clients were very satisfied with the solution, but were limited in the amount of data which could be processed. Jupiter is a joint SAP and IBM project intended to vastly expand the capability of the BI Accelerator solution and show that large databases will scale and perform on IBM infrastructure with the BI Accelerator solution.

Pronunciation notes
BI is pronounced B I
BIA is pronounced B I A
GPFS is pronounced G P F S
GUI is pronounced gooey
KPI is pronounced K P I
SAP is pronounced S A P
TB is pronounced terabyte
TOPAS is pronounced Toe pass
TREX is pronounced T REX

1a Time: 00:00 40sec (Title)
SAP's Business Intelligence Accelerator (or BIA) has been providing customers with super fast query processing to meet their business intelligence needs. Clients were very satisfied with the solution, but were limited in the amount of data which could be processed, with 28 blades representing the maximum BI Accelerator blade configuration. In 2007, SAP and IBM embarked on the Jupiter project, whose main intent was to vastly expand the capability of the BI Accelerator solution and show that large databases would scale and perform on the IBM platform with the BI Accelerator solution.

2a Time: 00:40 20sec (Agenda)
This demo is designed to quickly introduce the BI Accelerator solution and discuss the details of the Jupiter project jointly run by SAP and IBM.
We'll review the system configuration and environment,
our data model,
the test specifics,
and of course our outstanding results.

3a Time: 01:00 55sec (What is Jupiter?)
Jupiter is really the code name we used on the project, but it did have some meaning in that we were "Jupiter" sizing the databases we wanted to run in the SAP/BI Accelerator environment.

The project was designed to scale a DB2 database from 5TB to 25TB, with the 25TB database size being well beyond any BI Accelerator environment that was currently in existence.

The intent was to prove that BI Accelerator would be scalable with increasing data volumes and the increasing resources required to handle those volumes.

In order to ensure an objective review of our results, the Jupiter Project results were independently audited by WinterCorp, a well known organization specializing in Business Intelligence.

4a Time: 01:55 55sec (Project Scope)
It has already been mentioned that the BI DB2 databases used were from 5TB to 25TB. The project actually ran the suite of tests with 5TB, 15TB, and 25TB data volumes.

Two of the major tests performed and measured at each database size were
the BI Accelerator index creation and
the multi-user query loads.

The BI Accelerator index creation measured the time required to load the blade environment with data. The multi-user query tests would simulate activity for 100 to 800 parallel threads performing queries. BI Accelerator index creation and the multi-user query load tests were performed at each data volume size with varying blade configurations. For the Jupiter project, 27, 81, and 135 blade landscapes were utilized.

5a Time: 02:50 60sec (Architecture)
Now, let's talk about the SAP Business Intelligence Accelerator Architecture.

On the left hand side of the chart is a very common BI infrastructure as it exists today without the BIA. Data is placed in InfoCubes, and the BI analytical engine retrieves the information from the InfoCubes to perform its queries.

On the right hand side of the chart is the BI Accelerator. Now instead of explicit tuning by creating aggregates for each InfoCube, data is read to and stored in BIA blade memory and IBM's General Purpose File System (or GPFS) in a highly compressed, optimized, column-based representation of the data.

This process is referred to as the BIA index creation. With the large amount of blade memory and processing capacity, a high degree of parallelism can be realized and applied to query processing. Also, BIA can be integrated seamlessly and transparently into an existing BI solution.

6a Time: 03:50 60sec (Environment)

The Jupiter environment consisted of a database server, an application server and a BI Accelerator.

The database server consisted of an IBM System z9 S38 processor with three IBM DS8300 storage systems.

The application server was an IBM System p595 processor with one IBM DS8300 storage system.

For the BIA, we configured ten IBM blade centers with HS21 blade technology, ten IBM 3650 GPFS storage nodes, and two IBM DS8300 storage systems.

The test drivers resided on five IBM System x 335 processors.

The environment was substantial, as the intent was not to have any component introduce a constraint to our scalability and performance testing.

7a Time: 04:50 45sec (Data Model)
The data model for the Jupiter project included 78 InfoCubes and 30 billion, yes that was billion, records.

All InfoCubes were loaded into the BI Accelerator and all InfoCubes were accessed by the reports. The data in the InfoCubes was based on real life, customer defined scenarios along with additional benchmark scenarios. All query testing utilized standard BI functionality with no fine tuning.

Now, let's look at some of the screen shots during the InfoCube load. Note that this data loading process takes place on the System z9 with DB2 for z/OS. All actual screen captures in this demo were from loading and testing the 25TB database with the 135 blade landscape.

8a Time: 05:35 15sec (Process Chain screen)
The "SAP GUI Process Chain Display Planning View" lists the jobs for the InfoCube load process. These jobs populate the contents for Billing, Delivery, Sales Order, and the benchmark cubes in the 25 TB database.

9a Time: 05:50 20sec (Data Warehousing Workbench screen)
To monitor the InfoCube load processes, we use the "Data Warehouse Workbench Modeling display". Here we see 10 Billing cubes, 10 Delivery cubes, 10 Sales Orders cubes, and 48 benchmark cubes. This represents the full complement of the 78 cubes.

9b Time: 06:10 20sec (Data Warehousing Workbench screen – click on Manage)
Let's check the progress of one customer cube in particular – Fiscal 2010 Sales Order.

The yellow request indicator shows this cube is in the load process. When the load process request completes for this cube, its status will change to green, like the others below it.

10a Time: 06:30 35sec (Index Creation)
The first test suite was the BI Accelerator index creation. BI Accelerator indexes were created for all 78 InfoCubes. To achieve the best performance, various combinations of the number of InfoCubes processed in parallel and the number of partitions in which to split the fact tables were used.

The key performance indicators or KPIs for the BI Accelerator index creation were: the total data load time into the BIA, the load throughput per hour, and CPU utilizations on the BI server and the blade servers.

11a Time: 07:05 30sec (Global Work Process Overview - SM66 screen)
One of the tools we used is the Global Work Process Overview from the SM66 transaction. It shows the state of all active Batch and Dialog processes for the SAP BI application server. BI Index creation is the process where data is read from a row-based database and transformed into in-memory column-based files in a cluster of blades.

Let's look at another monitoring tool…

12a Time: 07:35 05sec (Simple Job Selection – SM37 screen)
This display is the Job Overview from the SM37 transaction.

12b Time: 07:40 25sec (Job Overview – SM37 screen)
SM37 displays the details for each background job. It shows the job status; start and end times; duration; and number of active, finished, and cancelled jobs at any given time.

The bottom of the display in the shaded area represents 10 cubes being indexed in parallel.

13a Time: 08:05 16sec (TOPAS screen)
The TOPAS monitor shows us the resource usage of the System p5. As you can see, the CPU utilization during the BI index creation phase was approximately 90%, which is obtained by adding kernel plus user percentages.

14a Time: 08:21 15sec (Web-based z RMF screen)
To monitor System z CPU, memory, and channel utilization, we used the web-based z RMF monitor. The current display represents the activity on the System z during the BIA index creation phase.

15a Time: 08:36 40sec (TREX ADMIN tool for blade monitor screen)
We have now transitioned to a BI blade monitoring tool. With the SAP TREX Admin screen we can get a snapshot of the index creation activity taking place on the blade servers. Each row of

the display provides statistics for a single blade server. Note that the CPUs are not overcommitted and that memory utilization is below the recommended 50% threshold, which is represented by the red line in the memory column.

Our configuration consists of 10 IBM BladeCenter chassis. Each chassis contains 14 blade servers so we are running with a total of 140 blade servers.

16a Time: 09:16 07sec (Back to Job Overview - SM37 screen)
Let's go back and take a look at one of the completed index creation jobs using the SM37 Job Overview transaction display.

16b Time: 09:23 25sec (Job log entries screen)
We are viewing the completed job for one of the Sales Order cubes. Each row of the display represents a cumulative load of records. The Job Log entries for this job show us that 486 million records were indexed into the BI accelerator.

17a Time: 09:48 22sec (TREX ADMIN at blade side screen)
Now let's take a look at the cumulative statistics for the index creation phase on the blade servers. The SAP TREX Admin screen shows us the CPU activity in red and memory utilization in green across all 135 blades during index creation. Each line represents a single blade.

18a Time: 10:10 32sec (Multi-User Reporting Tests)
At this point in time, index creation is complete, and data is ready for query processing. The second test suite is the multi-user reporting tests.

This test suite was designed to provide a realistic multi-user simulation on the BIA, with varying reporting capability and ensuring all data within the BI Accelerator was accessed.

The load drivers were used to replicate a multi-user query load and also provide loads from 100 to 800 parallel threads.

18b Time: 10:42 16sec (Multi-User Reporting Tests KPIs)
The key performance indicators for the multi-user reporting were:

reports and query throughput,

selectivity,

response times,

and CPU utilizations on the BI server and the blade servers.

19a Time: 10:58 32sec (Load Runner screen showing multiple user query)

With the index creation complete, let's look at a 25 terabyte 135 blade multi-user query run from the workload driver tool. As you can see in the upper shaded area, the workload driver increases the load to 100 threads. Each of these threads represents multiple users. The highlighted area in the lower portion of the display shows the response times for the query transactions.

20a Time: 11:30 26sec (Global Work Process Overview - SM66 screen)
The SAP Global Work Process Overview shows the state of all active processes running on the System p5 during the multi-user query test. The System p BI application server takes query requests from the workload driver and forwards them on to the BI accelerator for processing. The result is returned to the app server for rendering with final results sent back to requestor.

21a Time: 11:56 18sec (TOPAS screen)
Switching to the TOPAS display, we see that the System p5 CPU activity was very high during the multi-query run. This is due to resources needed by all the active processes displayed on the previous screen.

22a Time: 12:14 23sec (TREX Admin tool screen)
Let's look again at the blade server instantaneous activity using the TREX Admin tool, this time showing our query processing activity on the blades. Once again, we are primarily looking at the CPU and memory utilization.

22b Time: 12:37 21sec (TREX ADMIN at blade side screen)
Now let's take a look again at cumulative statistics for the blade servers. In the beginning we see the blade utilization during the index creation phase.

At the end of the display we see the blade utilization for four multi-user query runs, with the fourth still in progress.

23a Time: 12:58 38sec (Linear Scalability)
The scalability results literally speak for themselves.

As the database sizes increased and appropriate systems resources were increased, SAP/BI Accelerator on IBM infrastructure provided true linear scalability.

(start displaying boxes)
Index creation throughput, multi-user reporting throughput, and average report response time were consistent at each database size.

(all boxes displayed)
Such scalability was maintained as the number of records per report increased with each larger sized database.

24a Time: 13:36 37sec (Summary)

Project Jupiter results have been characterized as "very impressive" by BI experts. Today's BI Accelerator clients, or perspective clients, now have a highly scalable solution for their BI needs of today and into the future. No scalability limits were identified by the project. Project Jupiter has proven that SAP/BI Accelerator on IBM infrastructure can scale well beyond current solution limitations and provide true linear scalability to 25 TBs and beyond.

24b Time: 14:13 15sec (Closing Statement)

IBM's Infrastructure, together with SAP BI Accelerator are ready today to handle data warehousing requirements of tomorrow.