The need of BI Accelerator and its Complete Architecture

Applies to:

SAP NetWeaver 7.0 BI 7.0 with SP4 or higher

Summary

This article addresses one of the best BI solution's for Performance optimization and Lower TCO, The **BI Accelerator**, its need and architecture in a very simple way.

Author: K.Madhavi Latha
Company: HCL Technologies
Created on: 04th July 2007

Author Bio

Madhavi Latha is working as a SAP NetWeaver BI consultant in HCL Technologies.

© 2007 SAP AG

Table of Contents

© 2007 SAP AG

© 2007 SAP AG

Why SAP BI Accelerator?

Business Intelligence solutions have become a mission critical part of today's SAP landscapes due to these factors:

- Data explosion
- Increasing Users/analytics
- Frequent need for real time data
- Speed and flexibility to access actionable information (Performance issues)

These factors in turn have a negative effect on the design, setup and maintenance of BI solution's and became a challenging task.
Keeping in view these challenges SAP has come up with this new solution SAP NetWeaver BI Accelerator, This helps:
To expand the use of its Analytic server technology.
The customers to achieve a scalable, flexible and high performance solution.
BI Accelerator for High Performance BI is a performance Supercharger which does miracles for the apt Business environments.

Available performance optimizers

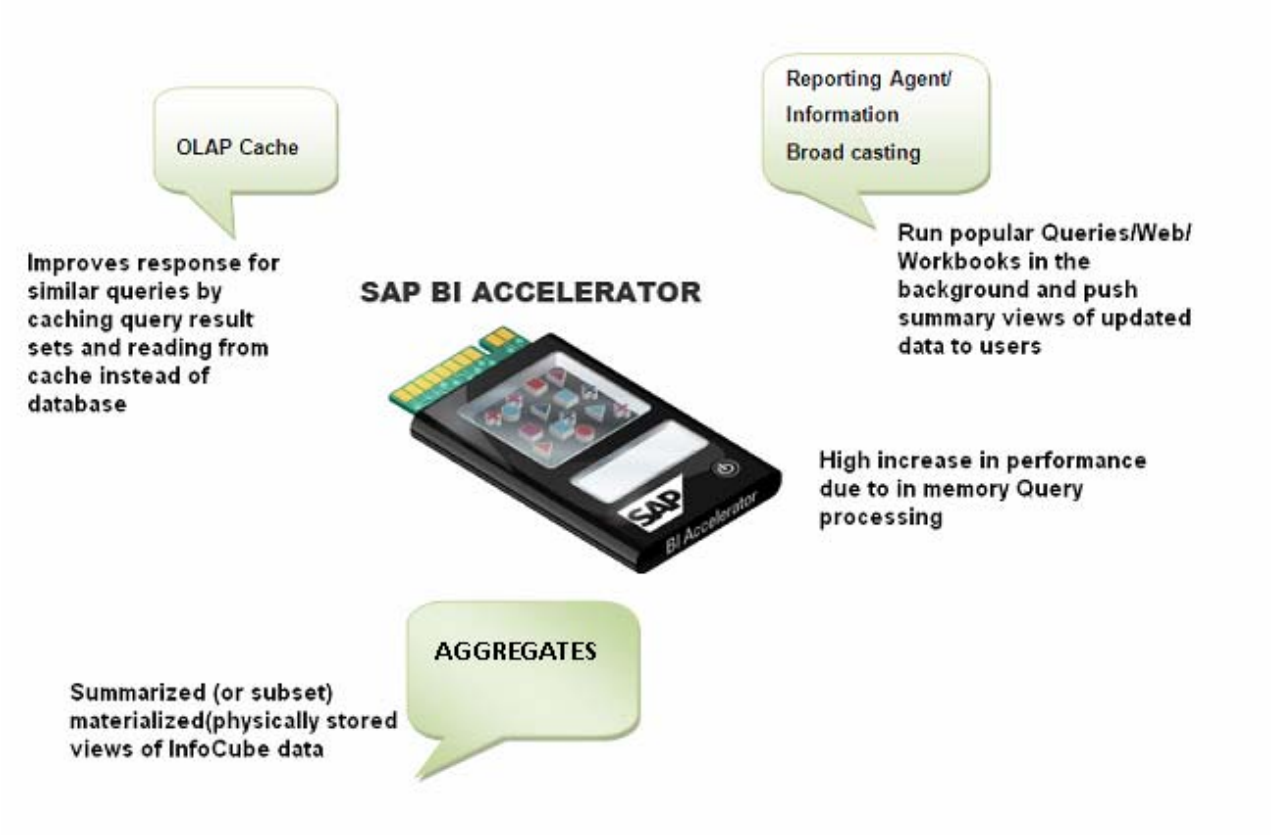

© 2007 SAP AG

BIA Architecture

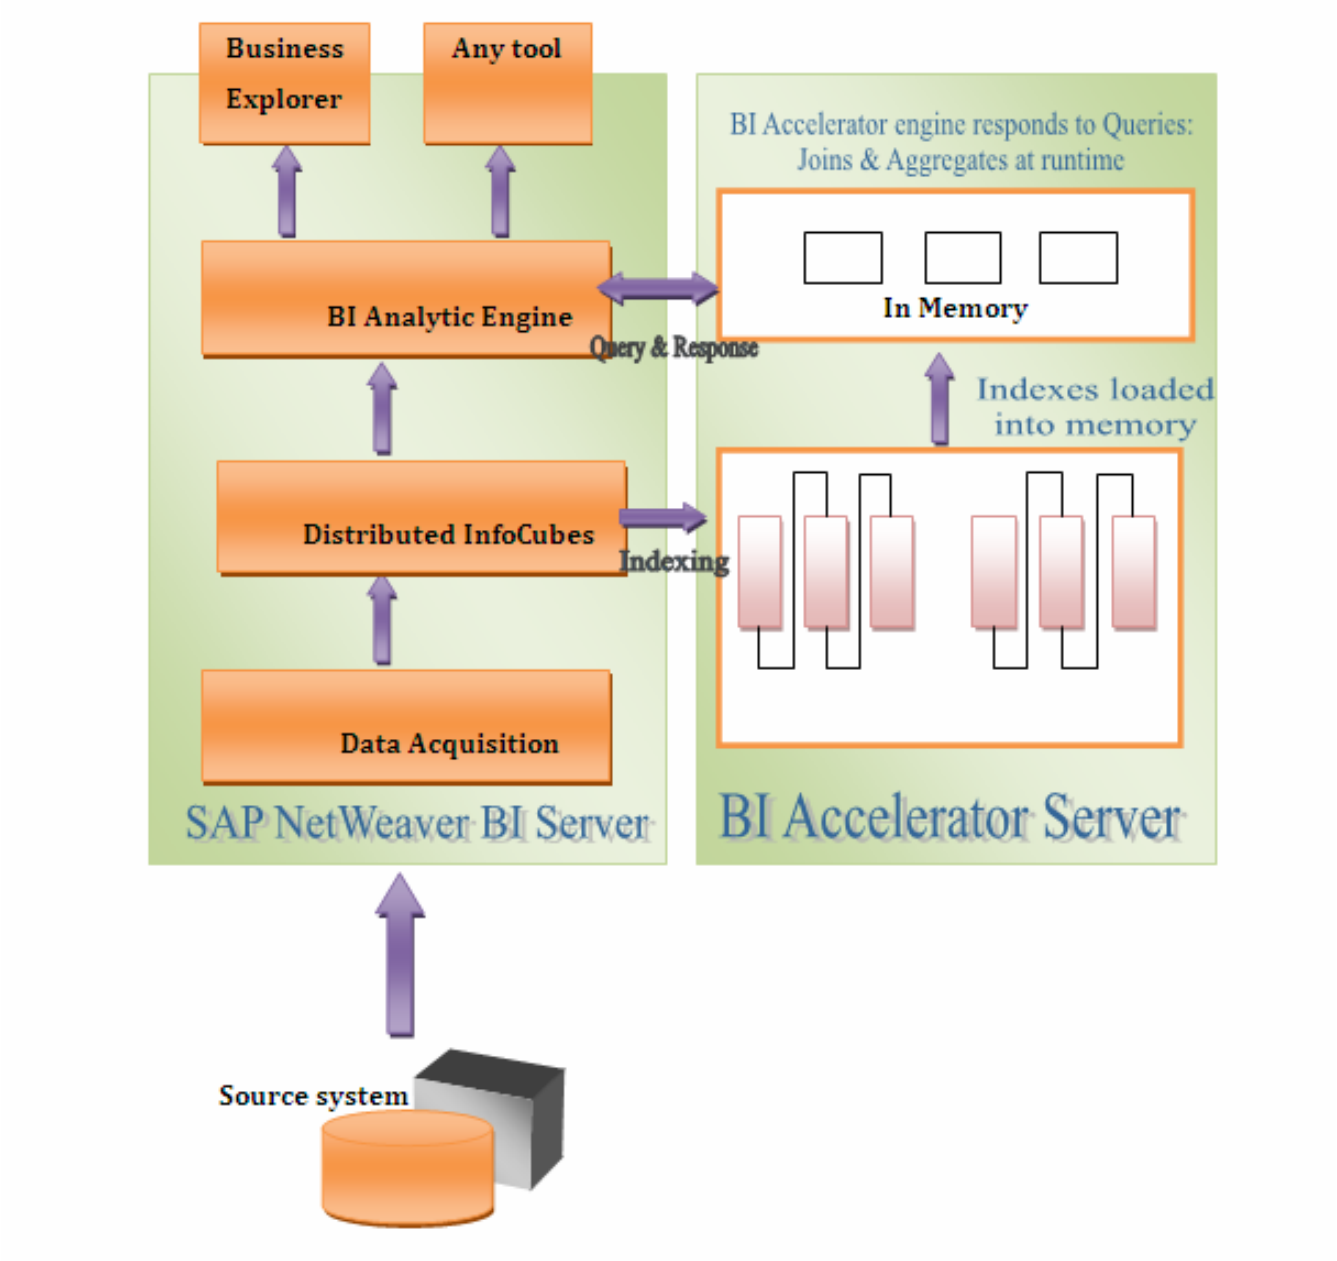

The BI Accelerator Server is a prepackaged hardware and software appliance that leverages SAP TREX search technology, combined with hardware technology developed by Intel Xeon 64-bit CPU blade servers and the blade servers must run on LINUX SLES 9.

© 2007 SAP AG

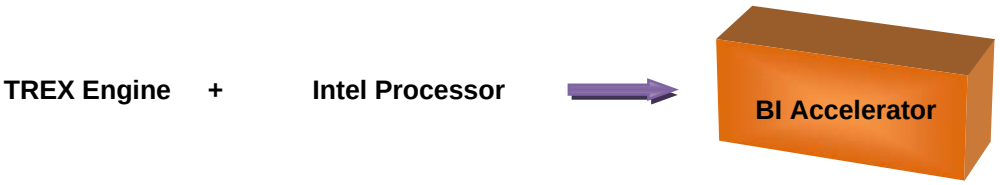

While TREX engine is meant for search and classification of unstructured data, BI Accelerator is used to handle structured data.
BI Accelerator engine and TREX are two different installations.
BI Accelerator engine is part of the Analytic engine that manages the BIA index.
The software allows to read, add or change 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.

Architectural features of BI Accelerator

The main Architectural features of the BI Accelerator are as follows:

- Scalable Multiserver Architecture
- In Memory Processing
- Optimized data structures
- Data compression using integers

BI Accelerator Query Architecture

After connecting BI Accelerator and BI instance through RFC, The BI Accelerator work flow is as follows:
Data from an SAP InfoCube is loaded into BI Accelerator
An index is created to this InfoCube 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.

© 2007 SAP AG

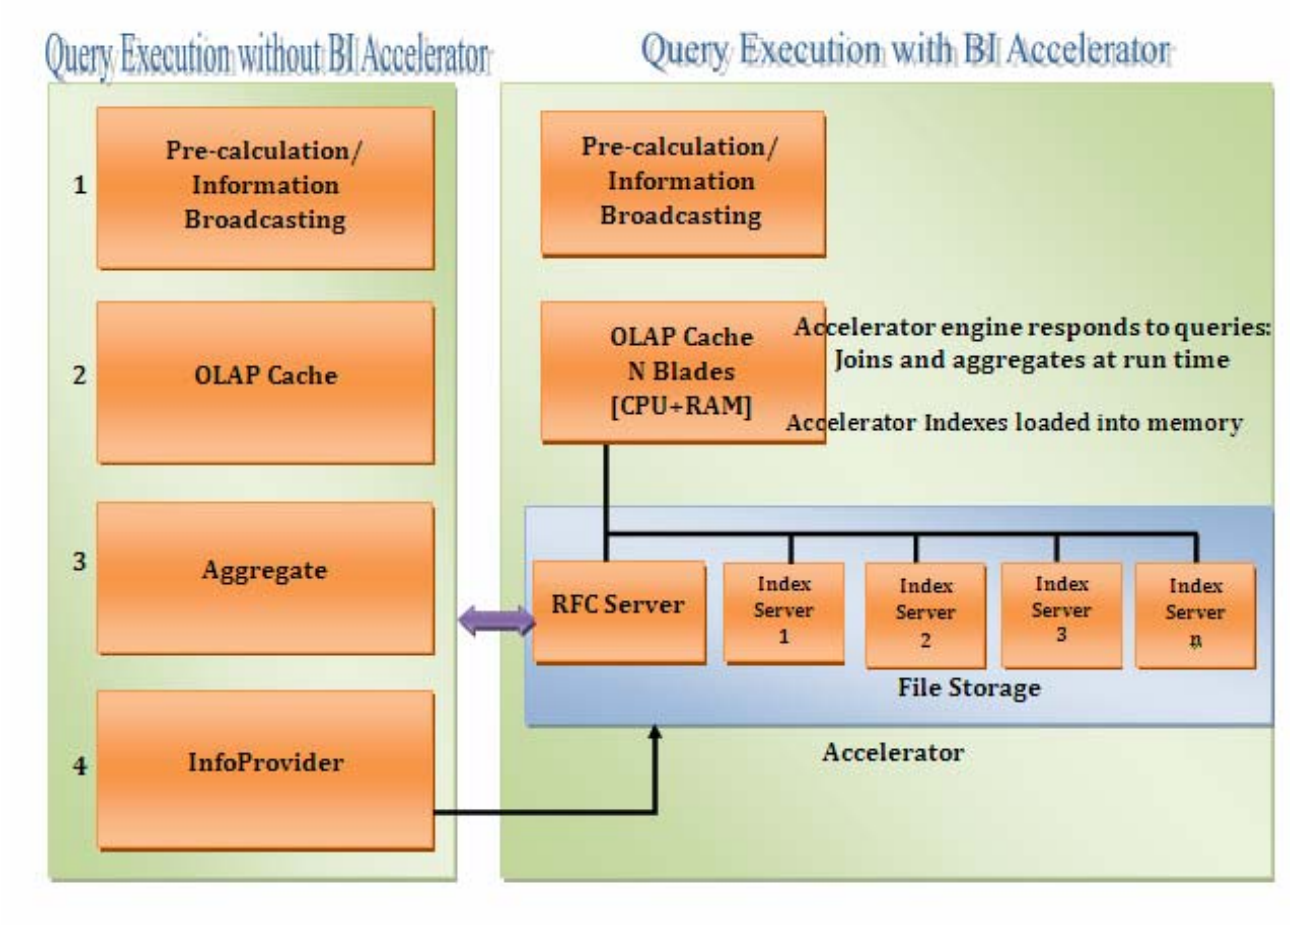
Query Execution without BI Accelerator
Query Execution with BI Accelerator
1
Pre-calculation/ Information Broadcasting
2
OLAP Cache
3
Aggregate
4
InfoProvider
Pre-calculation/ Information Broadcasting
OLAP Cache N Blades [CPU+RAM]
Accelerator engine responds to queries: Joins and aggregates at run time
Accelerator Indexes loaded into memory
RFC Server
Index Server 1
Index Server 2
Index Server 3
Index Server n
File Storage
Accelerator

© 2007 SAP AG

BIA Query Performance Boost-up factors

Vertical Decomposition

The Accelerator engine holds table data vertically, into columns that are stored separately. This makes more efficient use of memory space than row based storage (which happens in a conventional database), since the engine needs to load only the relevant attributes or characteristics data into memory.

Column storage

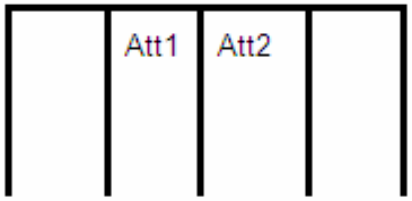

To find all Instances of an attribute value

1. Go to attributes column
2. Read its row values

Classical DB storage

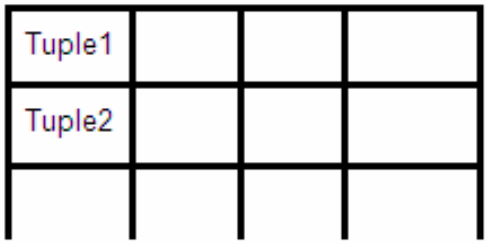

To find all instances of an attribute value:

1. Go to the first row
2. Check the attribute value
3. Go to the next row
4. Check the attribute value
5. Repeat for each row in the table

Horizontal Partitioning

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.

© 2007 SAP AG

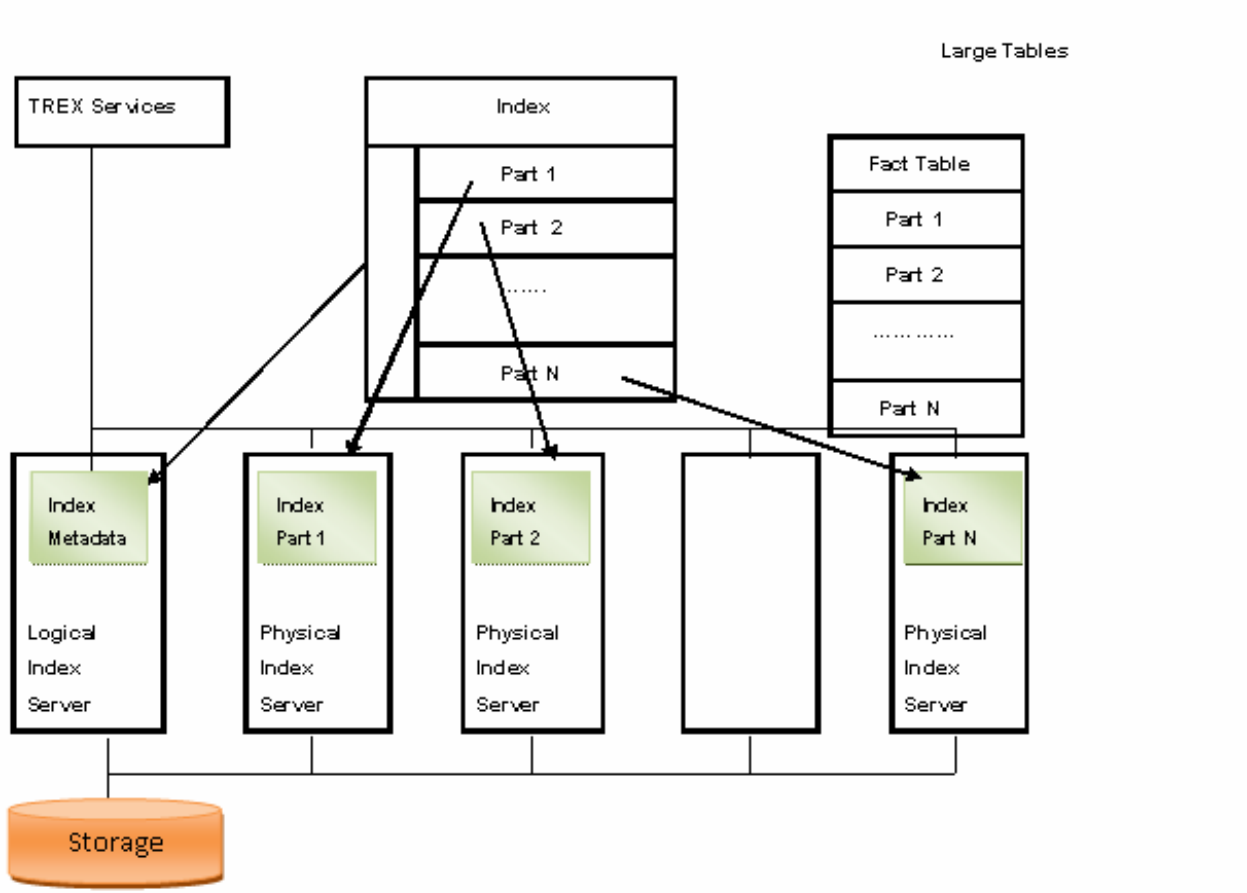

Smart Compression

BI Accelerator data is stored in read-optimized format. The data is compressed using integer coding and dictionary lookup.
Integer compression greatly reduces the average volumes of processed 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 10. Hence all memory processing can be done in main memory.
Integers represent the text or other values in tab cells, and the dictionaries are used to replace integers by their values during post-processing.
For example if a column has thousand rows and some of the cells contain long texts, efficiency is significantly increased by using a ten-bit binary number to identify the texts during processing and a dictionary to call them again afterwards. The datasets that have to be transferred and temporarily stored during the different processing steps are reduced on average by a factor of 10. This means that you can perform the entire query processing in the main memory and reduce network traffic between separate landscapes.

© 2007 SAP AG

TREX Storage for a
Single valued Attribute

Dictionaries

List of all used values
Values stored are compressed

Attribute Tables

Minimal number of bits used to represent values

Indexes

Compacted using sophisticated compression coding

Indexes for all Star schema tables

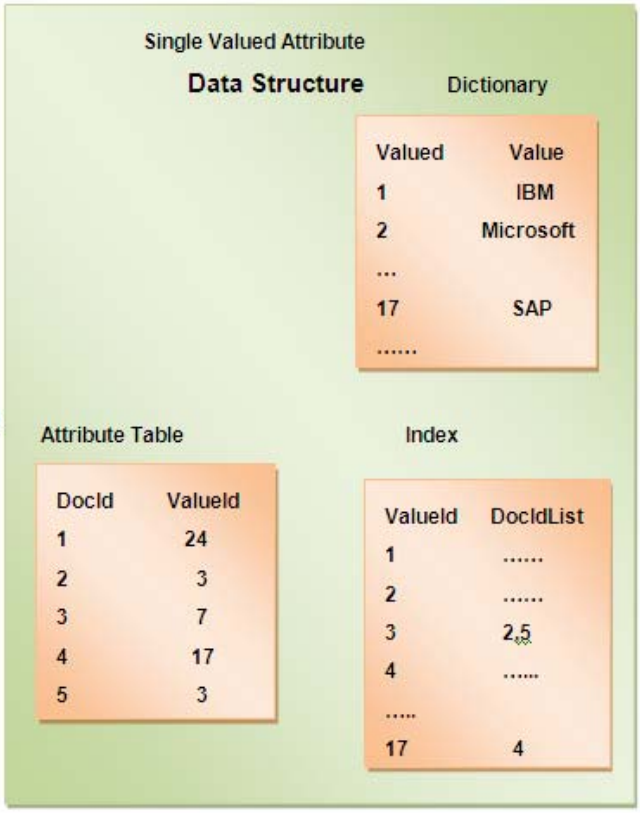

© 2007 SAP AG

BIA Rationale

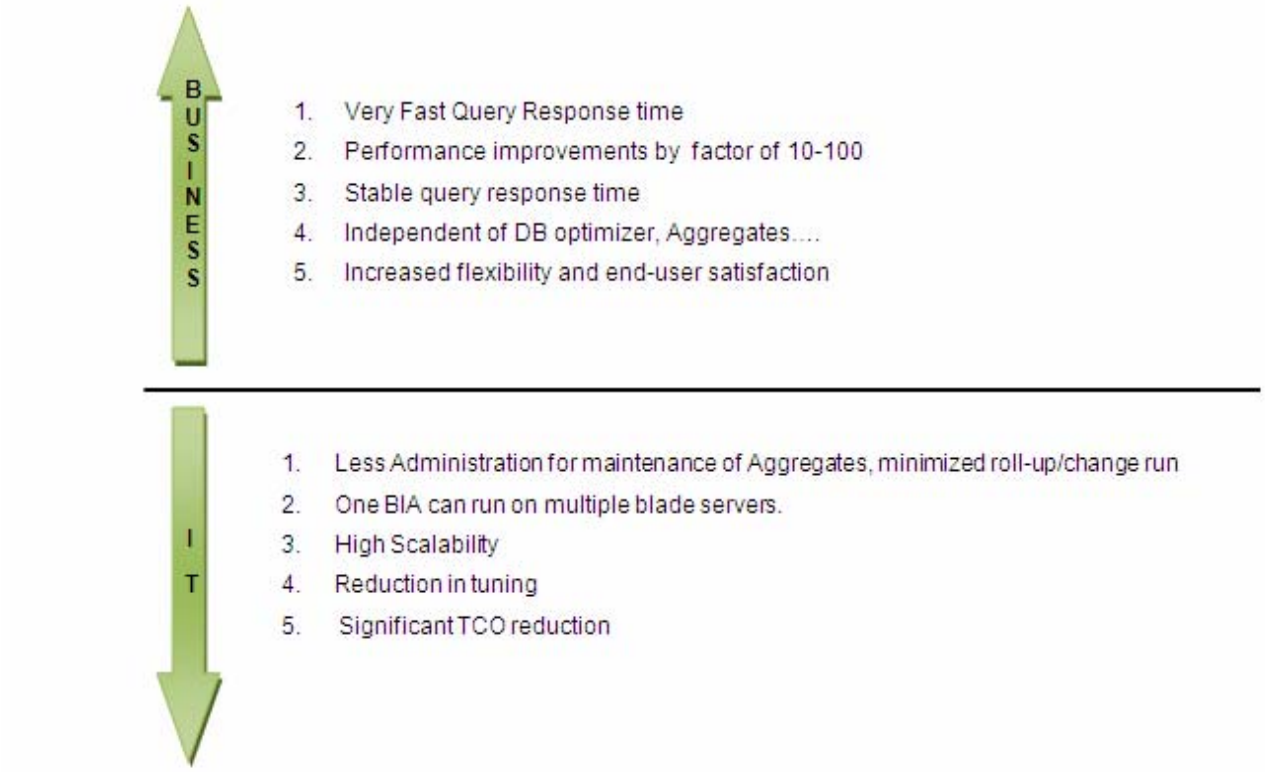

© 2007 SAP AG

Related Content

- https://www.sdn.sap.com/irj/sdn/go/portal/prtroot/docs/library/uuid/760410c5-0e01-0010-9984-ef6af105e902
- https://www.sdn.sap.com/irj/sdn/go/portal/prtroot/docs/library/uuid/3604c604-0901-0010-f0aa-b37378495537

© 2007 SAP AG

Copyright

© Copyright 2007 SAP AG. All rights reserved.
No part of this publication may be reproduced or transmitted in any form or for any purpose without the express permission of SAP AG. The information contained herein may be changed without prior notice.
Some software products marketed by SAP AG and its distributors contain proprietary software components of other software vendors.
Microsoft, Windows, Outlook, and PowerPoint are registered trademarks of Microsoft Corporation.
IBM, DB2, DB2 Universal Database, OS/2, Parallel Sysplex, MVS/ESA, AIX, S/390, AS/400, OS/390, OS/400, iSeries, pSeries, xSeries, zSeries, z/OS, AFP, Intelligent Miner, WebSphere, Netfinity, Tivoli, Informix, i5/OS, POWER, POWER5, OpenPower and PowerPC are trademarks or registered trademarks of IBM Corporation.
Adobe, the Adobe logo, Acrobat, PostScript, and Reader are either trademarks or registered trademarks of Adobe Systems Incorporated in the United States and/or other countries.
Oracle is a registered trademark of Oracle Corporation.
UNIX, X/Open, OSF/1, and Motif are registered trademarks of the Open Group.
Citrix, ICA, Program Neighborhood, MetaFrame, WinFrame, VideoFrame, and MultiWin are trademarks or registered trademarks of Citrix Systems, Inc.
HTML, XML, XHTML and W3C are trademarks or registered trademarks of W3C®, World Wide Web Consortium, Massachusetts Institute of Technology.
Java is a registered trademark of Sun Microsystems, Inc.
JavaScript is a registered trademark of Sun Microsystems, Inc., used under license for technology invented and implemented by Netscape.
MaxDB is a trademark of MySQL AB, Sweden.
SAP, R/3, mySAP, mySAP.com, xApps, xApp, SAP NetWeaver, and other SAP products and services mentioned herein as well as their respective logos are trademarks or registered trademarks of SAP AG in Germany and in several other countries all over the world. All other product and service names mentioned are the trademarks of their respective companies. Data contained in this document serves informational purposes only. National product specifications may vary.
These materials are subject to change without notice. These materials are provided by SAP AG and its affiliated companies ("SAP Group") for informational purposes only, without representation or warranty of any kind, and SAP Group shall not be liable for errors or omissions with respect to the materials. The only warranties for SAP Group products and services are those that are set forth in the express warranty statements accompanying such products and services, if any. Nothing herein should be construed as constituting an additional warranty.
These materials are provided “as is” without a warranty of any kind, either express or implied, including but not limited to, the implied warranties of merchantability, fitness for a particular purpose, or non-infringement. SAP shall not be liable for damages of any kind including without limitation direct, special, indirect, or consequential damages that may result from the use of these materials.
SAP does not warrant the accuracy or completeness of the information, text, graphics, links or other items contained within these materials. SAP has no control over the information that you may access through the use of hot links contained in these materials and does not endorse your use of third party web pages nor provide any warranty whatsoever relating to third party web pages.
Any software coding and/or code lines/strings (“Code”) included in this documentation are only examples and are not intended to be used in a productive system environment. The Code is only intended better explain and visualize the syntax and phrasing rules of certain coding. SAP does not warrant the correctness and completeness of the Code given herein, and SAP shall not be liable for errors or damages caused by the usage of the Code, except if such damages were caused by SAP intentionally or grossly negligent.