



How To... Plan for, Implement and Maintain SAP NetWeaver BW Accelerator Use Cases

Applicable Releases:

SAP NetWeaver 7.0 and later

IT Practice:

Business Information Management

IT Scenario:

Enterprise Data Warehousing

Version 1.0

February 2009

© Copyright 2009 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.

SAP NetWeaver "How-to" Guides are intended to simplify the product implementation. While specific product features and procedures typically are explained in a practical business context, it is not implied that those features and procedures are the only approach in solving a specific business problem using SAP NetWeaver. Should you wish to receive additional information, clarification or support, please refer to SAP Consulting.

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.

Disclaimer

Some components of this product are based on Java™. Any code change in these components may cause unpredictable and severe malfunctions and is therefore expressly prohibited, as is any decompilation of these components.

Any Java™ Source Code delivered with this product is only to be used by SAP's Support Services and may not be modified or altered in any way.

Document History

Document Version	Description
1.00	First official release of this guide

Typographic Conventions

Type Style	Description
<i>Example Text</i>	Words or characters quoted from the screen. These include field names, screen titles, pushbuttons labels, menu names, menu paths, and menu options. Cross-references to other documentation
Example text	Emphasized words or phrases in body text, graphic titles, and table titles
Example text	File and directory names and their paths, messages, names of variables and parameters, source text, and names of installation, upgrade and database tools.
Example text	User entry texts. These are words or characters that you enter in the system exactly as they appear in the documentation.
< Example text >	Variable user entry. Angle brackets indicate that you replace these words and characters with appropriate entries to make entries in the system.
EXAMPLE TEXT	Keys on the keyboard, for example, F2 or ENTER.

Icons

Icon	Description
	Caution
	Note or Important
	Example
	Recommendation or Tip

Table of Contents

- 1. BUSINESS SCENARIO 1**
 - 1.1 Enhanced Productivity 1
 - 1.2 Reduced Total Cost of Ownership..... 1
- 2. BACKGROUND INFORMATION AND PREREQUISITES..... 2**
- 3. SAP NETWEAVER BW ACCELERATOR PREPARATION AND SIZING FOR GO-LIVE.. 2**
 - 3.1 Preparation 2
 - 3.2 Sizing 3
- 4. UNDERSTANDING THE CAPABILITIES OF SAP NETWEAVER BW ACCELERATOR .. 3**
 - 4.1 Query Runtime..... 3
 - 4.2 Setting and Managing Expectations 4
- 5. QUERY AND INFOCUBE IDENTIFICATION FOR SAP NETWEAVER BW ACCELERATOR 5**
- 6. DATA MODELING AND DESIGN TIPS..... 6**
 - 6.1 Datastore Objects in SAP NetWeaver BW 7.0..... 7
 - 6.2 Lifecycle Management..... 8
- 7. FRONT-END TIPS AND TRICKS 8**
 - 7.1 Monitoring Progress 8
 - 7.2 Detailed InfoProviders 8
 - 7.3 Alternatives to OLAP Intensive Operations 9
 - 7.4 Understand the Approach for the Rollout – Have a Plan 9
 - 7.5 Miscellaneous..... 9
- 8. SAP NETWEAVER BW ACCELERATOR ADMINISTRATION 10**
 - 8.1 Maintenance 10
 - 8.2 Change Management 10
 - 8.3 Network Requirements 10
 - 8.4 Accelerator Index Maintenance 11
 - 8.5 Compression..... 11
 - 8.6 Database Indexes..... 11
 - 8.7 System Management..... 12
- 9. WORKING WITH LARGE DATA VOLUMES 12**
 - 9.1 FEMS Compression 12
 - 9.2 Package-Wise Data Read 13
 - 9.3 Miscellaneous 15
- 10. OTHER CONSIDERATIONS 15**

10.1	Aggregates	15
10.2	OLAP Cache.....	17
10.3	Other Processes.....	17
11.	FUTURE OF SAP NETWEAVER BW ACCELERATOR - ROADMAP.....	17
11.1	SAP NetWeaver BW Accelerator 7.2	17
11.2	SAP BusinessObjects Accelerator	18
12.	SUMMARY	18

1. BUSINESS SCENARIO

Compared with typical response times for standard company SAP NetWeaver BW systems, in which InfoCubes have been supplemented with manually defined aggregates, analytic query response times with the accelerator are often tens or even hundreds of times better. Scenarios with challenging response time requirements are numerous. For example, a call center for handling customer calls needs to be able to extract relevant information from customer records within a few seconds. In such scenarios, the IT department faces demanding service level agreements, and high performance is not a luxury.

The SAP NetWeaver BW Accelerator offers two main benefits to the companies that deploy it alongside an SAP NetWeaver BW system:

- Enhanced productivity - The increased performance and other features of the accelerator enhance the productivity of BW system users.
- Reduced TCO - The total cost of ownership (TCO) of the BW system is reduced by introducing the accelerator compared with other strategies for boosting performance.

1.1 Enhanced Productivity

The productivity of SAP NetWeaver BW system users is enhanced with the accelerator for several reasons:

- Performance - Response times delivered by the BW system for queries are often much faster. Even for large InfoCubes with millions of rows, results can be returned in seconds.
- Stable Response Times - The response times for repeated executions of a query or for execution of similar queries remain stable even under high system load. The response times also vary predictably under changes in data design and query design.
- Scalability - The accelerator can easily be extended incrementally with additional hardware.
- Robustness - The accelerator blade hardware is reliable and requires much less routine administration than legacy architectures.
- Usability - The performance and stability features encourage users to explore their data. With the accelerator, they can drill down to the facts and explore their data in an ad hoc manner without fear of timeouts or over consumption of system resources.

1.2 Reduced Total Cost of Ownership

Moving to the other half of the value proposition, companies can expect to enjoy reduced TCO for their SAP NetWeaver BW systems by deploying the accelerator rather than other productivity boosters for several reasons:

- No Aggregate Maintenance - Without aggregates to maintain, fewer resources are needed for routine administration.
- System Consolidation - Dependencies between the database and reporting are relaxed, so data loads and reporting no longer need to be scheduled serially. Because the accelerator takes over the reporting load, system copies made to handle that load can be consolidated.
- Longer Hardware Usage - Investment in new database hardware can be delayed or avoided. Growth in reporting demand can be accommodated without adding load to the existing BW landscape.
- Ease of Reporting - Long and inflexible approval processes for reporting can be replaced by shorter and more open processes. Changes in reporting requirements can often be accommodated at no additional cost.
- Simplified Data Modeling - Because manual performance tuning is no longer as necessary as in the past, companies can afford to develop and release new analytic scenarios more quickly to their lines of business.

SAP provides a tool for calculating the benefits a company can expect by deploying the accelerator. The ROI (return on investment) calculator for the SAP NetWeaver BW Accelerator is an easy-to-use

tool to qualify and quantify the value of deployment by measuring the ROI that can be expected from introducing the accelerator into a landscape. With the right input data, the tool calculates the ROI in precise numerical terms. The calculator is built with SAP BusinessObjects Xcelsius as an intuitive standalone tool embedded in a Microsoft PowerPoint file that can be customized for use in larger presentations. It is available on request from SAP account executives and others.

This document will provide those involved in a SAP NetWeaver BW Accelerator project with advice for planning, implementing and maintaining the solution for multiple use cases. It provides lessons learned from multiple customer implementations as well as undocumented tips and tricks to highlight key guidelines for achieving optimal accelerator performance. It will lay out the BW Accelerator Use Cases for customers implementing the solution as well as provide a number of options to consider when implementing and supporting the environment. The details of the solution will be supported by specific examples displayed throughout the document.

2. BACKGROUND INFORMATION AND PREREQUISITES

The SAP NetWeaver BW Accelerator is a hardware and software based optimization method for increasing SAP NetWeaver BW query performance. It is shipped as an appliance and operates on specifically configured hardware and requires additional SAP licenses. Although it can only be delivered in configurations certified by SAP, there are no dependencies with database vendor or operating system of BW systems where the accelerator is connected.

As a prerequisite to implementing the Accelerator solution, customers must have ordered the hardware from an SAP certified hardware partner: Hewlett-Packard, Sun Microsystems, Fujitsu Siemens or IBM. In addition, SAP NetWeaver BW release 7.0 or beyond is required. As always, running the software with the latest release level is suggested.

3. SAP NETWEAVER BW ACCELERATOR PREPARATION AND SIZING FOR GO-LIVE

3.1 Preparation

Before going live with SAP NetWeaver BW Accelerator:

- Assess the suitability of InfoCubes for the solution (see section 5 - Query and InfoCube Identification for SAP NetWeaver BW Accelerator)
- Assess the size of hardware needed for Go-Live
- Consider how much the data will grow and what this means for sizing. Generally customers size the appliance for 1 to 2 years, given a rolling data model where unused InfoCubes are deleted from the appliance
- Allow ample time (assuming you are already live with SAP NetWeaver BW 7.0) to test and understand the implications of the system in a controlled environment (i.e. quality or test system)
- Thoroughly test indexing, query performance, operating system behavior, network and connectivity
- Consider connecting the future productive blade system to a quality/test system before Go-Live for testing purposes. This will help alleviate any unforeseen problems with hardware before Go-Live

3.2 Sizing

As a rule, sizing of blades should be based on 50% memory utilization. All data for the InfoCube tables is contained in indexes on the accelerator. These are then loaded into memory. For sizing, this means the amount of physical data loaded should be below 50% of the total amount of physical memory. For example, there are six blades with 16GB each. That is a total of 96GB. Therefore, the amount of physical data should not exceed half of that (48GB). It is important to understand this is a 'rule of thumb' for sizing. Based on experience, this has been the case with a majority of customer scenarios; however there are some rare cases where more memory is needed. Therefore, it is highly recommended that the resources are tested and monitored before Go-Live in a quality or test environment to ensure enough resources will be available for productive use.

If the accelerator becomes overloaded, performance will degrade as there won't be enough memory for building temporary data structures while processing BW query requests. The more blades, the more widely distributed the data becomes across those blades, effectively taking advantage of parallel processing. Because of parallelism, generally speaking, more blades mean faster queries. Also remember to plan for at least one backup blade. Resizing will be required when more data or InfoProviders are indexed or more users are running queries against the solution. For more information on how to assess sizing before Go-Live refer to SAP Note 917803 – Estimating the Memory Consumption of a BIA Index.

There is also a service delivered by SAP AGS to help aid in many of these aforementioned situations. See the heading "SAP SERVICES AND EXPERTISE-ON-DEMAND FOR SAP BI ACCELERATOR" under the link <https://www.sdn.sap.com/irj/sdn/bia>

In addition the SAP QuickSizer should be used to capture sizing information for SAP NetWeaver BW Accelerator, <http://service.sap.com/quicksizer> - Create project and select SAP NetWeaver BW. Most of the data captured is the same as for BW sizing. Therefore, if a BW sizing exercise has already been performed, copy & paste the values for the accelerator sizing. Note that a new QuickSizer project must be created, otherwise the accelerator options will not appear. If a BW system already exists (loaded with data), the program and method described in SAP Note 917803 (mentioned above) is recommended.

Also, refer to the thread below on SDN for additional information.

<https://forums.sdn.sap.com/thread.jspa?threadID=776332&tstart=0>

4. UNDERSTANDING THE CAPABILITIES OF SAP NETWEAVER BW ACCELERATOR

4.1 Query Runtime

The table below breaks out BW query runtime into four different buckets:

Figure 1: Query Runtime Break Out

Data Read Time:	Time spent reading and aggregating data from the database or the accelerator
OLAP Time:	Time spent performing calculations and conversions, displaying hierarchies, executing security checks, etc. These operations are done within BW's analytic engine
Network Time:	Time spent transferring the result set via the network to the client
Front-End Time:	Time spent returning data to the user and formatting the results

The accelerator selects data from its self-contained memory within the appliance which improves data read time, ultimately lowering query runtimes as follows:

- The most significant impact on performance is on read (and aggregation) time
- There is little to no impact on calculation (OLAP)
- There are no performance gains with regard to the network throughput

4.2 Setting and Managing Expectations

From the beginning, it is important to set the right expectations when implementing the solution. Most queries experience solid runtime performance increases. However, it is important to understand that the current release of SAP NetWeaver BW Accelerator does not improve the runtime for every query. The initial release is explicitly targeted at reducing data selection and aggregation time. There are scenarios where the accelerator does not boost query performance significantly compared to the traditional performance tuning measures (like aggregates), e.g. an aggregate created specifically for a single query. Therefore, position the accelerator based on its intent, e.g. queries with high database time not OLAP / front-end time. The table below can be used to help set expectations based on the current capability of the product.

Figure 2: Understanding the Capability of the Accelerator: Setting Expectations

<u>BW Accelerator HELPS Improve Performance If:</u>	<u>Little Improvement With The Accelerator If:</u>
<ol style="list-style-type: none"> 1. High database read time 2. Ad-hoc analysis or unpredictable query patterns 3. High ratio for number of records Selected versus records Transferred 4. Queries routinely needing many aggregates 	<ol style="list-style-type: none"> 1. Very fast queries without BW Accelerator 2. High network or front-end time 3. Complex authorization checks 4. Transferred amount of data > 500,000 (rule of thumb) 5. High OLAP/calculation time

As an added benefit, the accelerator also contributes to reducing the total cost of operations (e.g. aggregate design, data latency, performance tuning, database utilization, etc.) for virtually all scenarios. By using the solution, customers are able to remove many, if not all aggregates. By doing this, the cost to maintain and the corresponding storage costs are eliminated. There is on longer a need for Basis or performance resources to be permanently allocated to chase BW query performance issues. Because all queries using the accelerator are executed on the appliance, the database utilization significantly decreases, leaving other applications with the ability to more effectively use that processing time.

SAP NetWeaver BW Accelerator contributes to reducing data load time e.g. roll-up and change run times. This can increase the ability to meet daily data availability service level agreements. No longer having the uncertainty of knowing how long and when the change run will complete makes it much easier to predict when data will become available to users for reporting purposes. With queries relying on the accelerator and not aggregates, the change run should execute in roughly the same amount of time each time it is executed.

5. QUERY AND INFOCUBE IDENTIFICATION FOR SAP NETWEAVER BW ACCELERATOR

There are two schools of thought when determining which InfoCubes to put into BW Accelerator: The first and easiest is to put the entire environment into the appliance. This can appear costly up front, however the enterprise-wide performance gains, decreases in DASD usage, time removed from the nightly load window and reduction in resources chasing queries with dismal performance as well as application server and database server performance can significantly offset the upfront costs of the solution.

The other approach would be to individually identify those queries with high data read/data manager time and large aggregation e.g. high ratio of selected records in comparison to the number of records being transferred. To help identify queries best suited for the accelerator, choose from four possible tools per Figure 3 below.

Figure 3: Tools to Identify SAP NetWeaver BW Accelerator Query Candidates

TOOL	METHODS FOR ACCELERATOR QUERY IDENTIFICATION
1. BI Statistics Queries	Execute queries on business content delivered technical statistics InfoCubes to determine the query performance. e.g. query 0TCT_MC03/0TCT_MC03_Q0200
2. Workload Monitor	Use the Workload Monitor (transaction code ST03N) to review the statistics captured for any period of time
3. BI Statistics View	Review statistics in table RSDDSTAT_OLAP to identify those queries which have a high aggregation ratio. Select events 9000, 9010 and 9011 to obtain statistics for data manager time, number of selected records and number of transferred records
4. Query Monitor	Execute and debug queries in RSRT. Capture statistics of the query to identify queries with high Data Manager time. You can also see the time from other components (RFC).

For more information, refer to BIA BLOG Series, Part II: Preparing for a SAP NetWeaver BI Accelerator Implementation.

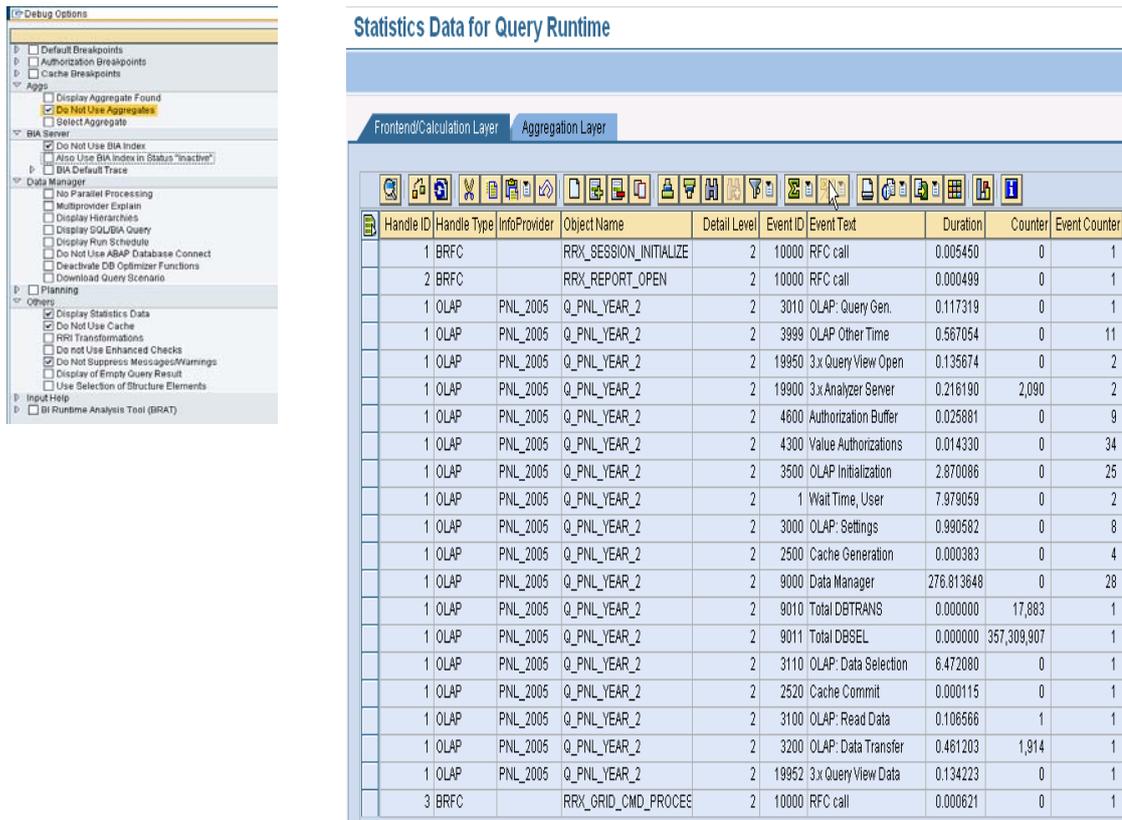
<https://www.sdn.sap.com/irj/scn/weblogs?blog=/pub/wlg/10009>

There is also a service delivered by SAP AGS to help determine cubes suitable for your accelerator solution. See the heading “SAP SERVICES AND EXPERTISE-ON-DEMAND FOR SAP BI ACCELERATOR” under the link:

<https://www.sdn.sap.com/irj/sdn/bia>

Figure 4 provides the statistics of a query using the Query Monitor (transaction RSRT). To effectively use this transaction, choose **Execute and Debug with Debug Options → Display Statistics Data → Do Not Use Cache → Do not Use BW Accelerator → Do Not Use Aggregates**. This query appears to be a good candidate for the accelerator because the read time – Data Manager is over 93% (276.81 out of 296.93 seconds). The accelerator solution should help alleviate most of the read time for this query.

Figure 4: Statistics Data for Query Runtime in the Query Monitor -RSRT



6. DATA MODELING AND DESIGN TIPS

The implementation of SAP NetWeaver BW Accelerator requires no changes to the existing best practice modeling guidelines. However, revisiting the design of the warehouse is an important first step to the implementation. It is important to understand the accelerator is not a replacement for poor performance due to a poorly designed data warehouse.

1. The accelerator is targeted at tuning end users' query performance
2. The accelerator will have a solid impact on the data warehouse with BW, specifically an accelerated change run as well as accelerated look ups from large objects when reading through the Data Manager and not directly from the database
3. Data warehouses which effectively perform and scale, rely heavily on well-defined architectures. This is a separate topic that needs special attention in the data warehouse but is not the focus of an accelerator deployment

There are customer scenarios where critical stakeholders of the BW system decide the accelerator is the only solution to a well-performing business analytical system. This is not always true and the entire BW landscape should be taken into consideration. To ensure any implementation is following best practices or recommended guidelines, refer to the BW Data Modeling and Front-end Design link below

<https://www.sdn.sap.com/irj/sdn/nw-bi?rid=/webcontent/uuid/c96f4c8d-0c01-0010-4d80-df184899323f>

Although the accelerator is a large piece of the BW puzzle, it is not the only one. There are also sizing considerations which affect both memory and CPU. Having a solid warehouse design will save money on the purchase of the appliance. What the accelerator does do is reduce the amount of ongoing performance-related activities. When accelerator performance makes manual performance tuning less

critical, companies can develop and release new analytical scenarios more quickly to their lines of business.

With the implementation of SAP NetWeaver BW 7.0 and the corresponding accelerator, a well-performing BW system can now be achieved with minimal effort and maintenance. See below the list of data warehouse activities and their associated solution:

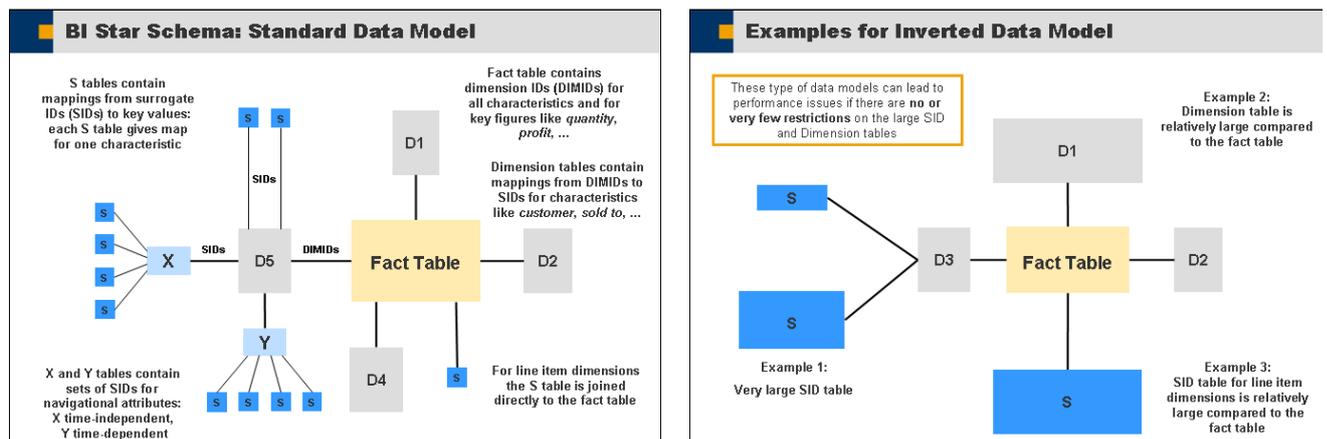
- Extract and stage data as fast as possible - achieved by write-optimized DSOs and use of DTP's
- Retrieve data as fast possible - use of the accelerator solution
- Select from massive amounts of data - use of the Accelerator solution
- Allow data to be indexed for reporting optimization in a short amount of time - use of the accelerator solution
- Archive old data but have readily available - use of Near-Line Storage

6.1 DataStore Objects in SAP NetWeaver BW 7.0

If applicable, use write-optimized DataStore Objects (DSO) as an enterprise data warehouse or inbound layer (as opposed to standard DSOs that are neither fully optimized for reporting or loading). Using write-optimized DSOs will remove the activation time from the data flow. Such design will also position the solution for greater use with the next release of BW where it is planned to enable the accelerator for DSOs. This release will provide the capability to load data from a DSO directly to the accelerator, without the need for persisting data in an InfoProvider. For more details see section 11 – Future of SAP NetWeaver BW Accelerator – Roadmap.

Additional considerations are needed when using the accelerator in conjunction with inverted data models, e.g. large master data tables (often larger than fact tables) as the joins are done at query runtime. If there are no filters in the queries, millions of SIDs and DIMIDs could potentially be joined which is typically expensive, regardless of whether the accelerator is used or not.

Figure 5: Standard vs. Inverted BI Star Schema Data Model



When working with inverted data models or generally large master data tables, consider the suggestions below:

- Avoid “unselective joins” by leveraging filters in queries
- Redesign the data model if possible to decrease the size of master data tables
- Consider making the dimension a line item dimension
- Use “flat” accelerator indexes as described in SAP Note 1047527 – “Flat” BIA Index: Indexing Performance With Line Items, or SAP Help http://help.sap.com/saphelp_nw04s/helpdata/en/16/d7a54216b6c611e1000000a1550b0/frameset.htm

- Delete unused master data
- Run "Analysis and Repair of BI Objects" (transaction RSRV) checks to clean up dimension data
- Leverage Dimension Function Index (DFI) for BW Accelerator when it becomes available. For more information, refer to SAP Note 1074559 – Dimension Function Index - Support in the BIA

6.2 Lifecycle Management

The principals of data lifecycle management still hold true with the implementation of the accelerator solution. In order to keep costs at an absolute minimum and performance at its highest, it is important to remember the principal of "Only access the data which is needed to report on". Very easily identifiable dimensions such as time or location (and others) should always be considered when designing the reporting layer. Consider using logical partitioning of the InfoProvider to ensure the accelerator has smaller subsets of data to work with. Although not required, logical partitioning will result in improved memory handling for the accelerator and although relatively minimal, performance can be increased. For example, consider the scenario with 12 logically partitioned InfoProviders by month, with only the two most recent months frequently queried. In this case, consider indexing only the two most frequently queried InfoProviders in the accelerator and create aggregates for the remaining 10. Keep in mind, however, that maintaining aggregates has a support cost associated and can affect the master data change run. By only loading the most frequently queried data into the accelerator, the appliance will have additional memory available to effectively use while processing query requests or to be allocated for future growth.

7. FRONT-END TIPS AND TRICKS

The accelerator is transparent to end users, requiring no changes to the existing queries. It is also recommended to follow front-end tuning and best practices for optimal query performance whether using the accelerator or not. To ensure front-end best practices are being followed, refer to the Front-end Design link below.

<https://www.sdn.sap.com/irj/sdn/nw-bi?rid=/webcontent/uuid/c96f4c8d-0c01-0010-4d80-df184899323f>

7.1 Monitoring Progress

Like any new project, after the implementation of the accelerator, it is important to monitor and analyze its performance e.g. RSRV performance tests and accelerator load. Performance activities should be focused on queries not achieving desired performance gains. Refer to Figure 3, Tools to Identify BW Accelerator Query Candidates, to further identify queries that do not meet performance expectations.

7.2 Detailed InfoProviders

Often times, InfoCubes are not effectively designed and contain very granular data. In many instances, an accelerator implementation will help with performance to an extent, but is limited to what it is capable of based on the data model design currently in place and size of the accelerator. Generally speaking, the more blades in the accelerator the better performance as the appliance relies heavily on parallelism by distributing the data across all blades. If large volumes of granular data exist and remodeling the InfoCube in a more effective manner is not a possibility, the queries must be designed so that a mandatory filtering to mitigate "runaway" queries which retrieve vast amounts of unneeded data is used.

7.3 Alternatives to OLAP Intensive Operations

As this is the first release of the solution, keep in mind its initial intention was not focused on OLAP operations. Section 11, Future of SAP NetWeaver BW Accelerator – Roadmap, will detail future enhancements in relation to OLAP. Until then, if performance in this area is not what was expected, redesign the query to try to avoid areas of functionality not currently part of this release of the accelerator; sorting, formula calculation before aggregation, exception aggregation, top-n or bottom-n, zero suppression in query properties, conditions-and-exceptions reporting, etc. (See Figure 2). If queries consist of numerous complex calculations, consider splitting them into several unique queries. Select the right read mode for your queries. For those queries involving hierarchies with many nodes, select "Read data during navigation and when expanding the hierarchy" option to avoid reading data for the hierarchy nodes that are not expanded. In regard to exception aggregation, please see SAP Note 1257455 - Exception Aggregation: Performance and Memory Requirements.

7.4 Understand the Approach for the Rollout – Have a Plan

Prioritize accelerator requirements for the implementation. Determine the approach for implementing the solution. Does it make the most sense to enable the accelerator over a period of time with the worst performing queries first or should it be enabled all at once? How long will the entire rollout take? As the first users experience the performance gains with the accelerator, other users will quickly demand access to it. Whichever method is chosen, have a plan and actively involve the users to determine which scenario is best given the circumstances. Regardless of the approach, ensure the appliance is sized accordingly.

7.5 Miscellaneous

When working with large query result sets, use multi-dimensional navigation to drill down to details rather than having a large initial report. The accelerator is not a solution for poor query or data model design.

For many scenarios, using Information Broadcasting to increase query speed (cache) or delivery (SAP Enterprise Portal or email) is still relevant even with the accelerator. This functionality should be considered on a case by case basis. Depending on the size of email generated, time it takes for query to execute, etc. the accelerator solution should still be considered.

If using Analysis Process Designer (APD) currently, consider disabling the accelerator for APD as these queries may be handling extremely large data volumes which could negatively impact performance. If this is not an issue, please see SAP Note 1231794 – APD Using BIA to Read InfoProvider Data.

As the accelerator may enable more efficient reporting on detailed InfoProviders, the report-report interface (RRI) may not need to be used as often, which will ultimately require less end-user training as well as remove complexity from the existing environment.

For additional query performance tips, refer to the white papers and blogs available on the SAP Service Marketplace or the SAP Developer's Network. A good first read is "Understand Query Performance in NetWeaver BI and BW Accelerator" by Thomas Zurek.

<https://www.sdn.sap.com/irj/sdn/weblogs?blog=/pub/wlg/5644>

8. SAP NETWEAVER BW ACCELERATOR ADMINISTRATION

First it is important to read the Accelerator's Technical Operations Manual. The information included in this document is important for all technical personnel who are involved with the accelerator.

<https://www.sdn.sap.com/irj/sdn/go/portal/prtroot/docs/library/uuid/46e11c10-0e01-0010-1182-b02db2e8bafb>

For optimal accelerator performance, ensure all of the prerequisites have been met.

8.1 Maintenance

Apply the latest SAP Notes relevant to the accelerator and implement the latest revision and BW Support Pack. For example, significant improvements have been made in memory management from Revisions 46-49. Upgrading to the most recent revision could increase your memory management by 30-40%.

- Fix for minLifeTime of temp indexes -> Memwatch
- Fix for overflow in bitvector
- Fix for memory issue if package-wise read is active
- Memory consumption of mergeIndex call
- Performance improvement for searches on line items
- Fixed memory leak if FEMS_COMPRESSION is used
- Reduced memory usage of TREXIndexServer during query execution

Load balancing, typically done by the Basis team, will help to more effectively manage all existing database and application servers, ultimately avoiding system bottlenecks during indexing and roll-ups of data. This activity can be performed as part of the set up of the appliance and should not need to be revisited until servers are added or removed from the landscape.

Minimize any accelerator activity not related to query execution during normal business hours e.g. reorganization and index creation/roll-up, as this will prevent any user disruption. Normally as these activities are occurring, an index may become unavailable which will result in a query request being sent to the database or an aggregate, which may result in suboptimal performance

8.2 Change Management

Although there is not a strict technical change management for indexing InfoCubes, it is important to design a process around adding new InfoCubes to the accelerator. First you should determine whether the InfoCube makes sense to index. Second, assess the estimated size of the InfoCube by running the program from SAP Note 917803 (mentioned in section 3). Third, assess whether the amount of available memory is sufficient to add the new InfoCube(s). It is also important to understand that normal data growth with the addition of newly added InfoCubes can cause serious problems with memory handling when sufficient resources are no longer available. See the sizing section for more information.

8.3 Network Requirements

A dedicated network connection with a minimal amount of jumps is required to ensure optimal performance. The network speed between the BW system and the accelerator must be 1 Gbit/sec or faster. It is also recommended that the accelerator blades be on the same subnet as the BW system and exclusively used by the accelerator. The reason for the strict requirements is extremely large packets of data could potentially be sent over the network. If a dedicated network connection is not available, there could potentially be a negative impact on other applications.

8.4 Accelerator Index Maintenance

In certain cases consider deleting and rebuilding the accelerator index for an InfoCube on a regular basis (e.g. once a month) to ensure a good reporting performance. There is no “hard” threshold as it depends on the amount of activity. Some customers choose to do this every month and others wait until there is noticeable performance deterioration. Monitoring the performance over a period of normal activity will be the best judge as to how often an index rebuild may need to take place. This will be different for every cube.

Once the frequency is determined the task can be automated in process chains with the use of process type **Initial Activation and Filling of BIA Indexes**. Rebuilding is especially important if deleting requests from an InfoCube and/or compressing the data in an InfoCube on a regular basis. If many requests are compressed or deleted over time, there will be a growing discrepancy between the number of records in the accelerator fact index and the InfoCube fact table which will ultimately impact query performance (query results are the same but the number of records are different between BW and the accelerator).

If a request is compressed, this action has no impact on the fact table index on the accelerator as there is only one index for the E and F fact tables in the accelerator. If a request is deleted from an InfoCube with an accelerator index, the data in the fact index is not deleted and instead, the ID of the deleted request is removed from the index of the package dimension. The deleted records from an InfoCube are removed entirely but are marked as hidden in the accelerator index. When marked as “hidden,” these records are not considered when creating the result set. Until the accelerator index is deleted and rebuilt, the number of records in the fact table and the index will not match which could lead to confusion.

Reorganizations of indexes should be executed or scheduled periodically. This will optimally redistribute the indexes across all accelerator blades. An alert will appear in the monitor, transaction RSDDBIAMON2, when the system believes a reorganization will benefit performance.

8.5 Compression

With the accelerator there is no difference between compressed and uncompressed (E and F) fact table handling. Thus, InfoProvider compression is only necessary for the InfoProviders with the following features:

- High number of cancellation requests (eliminate many records with 0 values)
- Non-cumulative key figures (inventory providers)
- Significant number of requests (for request maintenance performance)
- High number of partitions in the F table
- Real-time InfoCubes with slow B-tree database indexes (for example Oracle)
- Initial indexing (as requests might differ in size)

8.6 Database Indexes

Database indexes for the InfoProviders are less relevant once the accelerator is implemented. If considering deleting these indexes, see the scenarios below which still require secondary database indexes:

- Initial filling of accelerator indexes
- Open hub extractions
- Data marts for transferring data from an InfoCube to another InfoProvider
- InfoSet queries
- Queries using Request ID as a characteristic (typically an IT scenario, not a business requirement)
- When accelerator is deactivated for the query containing virtual InfoObjects

- In 3.x queries that still have calculated key figures with the "Calculate before Aggregation" property on.

8.7 System Management

If there are hundreds of temporary indexes in the accelerator, consider stopping and restarting all services on the appliance or deleting temporary indexes (transaction TREX ADMIN) to remove unused indexes from the accelerator memory. For more information, refer to SAP Note 1132962 – BIA 7.00: Deletion of More Than 1,000 Temporary Indexes.

The stability and performance of the accelerator server may be negatively impacted and queries may be terminated if it is overloaded. An accelerator overload may occur if too much data is loaded into the appliance, if the query/user load is so high temporary memory consumption becomes too large, or with Application Process Designer (APD) queries on large volume InfoCubes. Thus it is important to monitor and optimize the accelerator utilization as needed. In the case of an overload, several options are available:

- As data volumes and/or users grow, the appliance can easily scale with the growth by adding additional blades
- For temporary purposes, delete indexes with low priority if applicable and disable the accelerator for APD.
- Distribute data optimally across the accelerator blades. For more information, refer to SAP Note 1163149 – BIA 7.00: New Reorg Parameters
- For additional information about an accelerator overload, refer to SAP Notes 1093719 – BIA Capacity Utilization: Preventing Overloading, 1132572 – BIA Monitor: Adjusting Revision and Overloading Checks and 1223230 – Monitoring the BI Accelerator in Overload Situations.

9. WORKING WITH LARGE DATA VOLUMES

Additional performance considerations are required when requesting and transferring large data volumes from the accelerator. The first recommendation is to reevaluate whether those queries that are transferring these volumes are still relevant. Can the query or InfoCube be remodeled in a more effective manner? Scenarios with a high number of records transferred often result in high processing times within the analytic engine of BW and a high network load between the accelerator and the BW system. Possible causes include:

- Wide record length attributed to hidden key figures and free characteristics
- Queries containing all characteristics in the data model
- "Dumps" of the data without filter restrictions
- If large volumes exist, begin by reading SAP Note 1018798 – Reading Large Data Volumes From BI Accelerator

9.1 FEMS Compression

Consider activating FEMS compression if on revision 48 or higher. The database and SAP NetWeaver BW Accelerator can return data with different levels of detail. An important difference is the accelerator can return data that takes into account the structure elements. This means it can return information about the structure element (FEMS = Formula Element Selections) to which a certain value belongs. The database explodes the structure elements and returns single records that must be assigned to the correct structure elements by the OLAP. With SAP Notes 1074953 – FEMS, BIA Optimization and 1118425 – BIA 7.00: Activate Performance Feature FEMS Compression, the accelerator usually returns fewer records than the database. The amount of data read due to the convex hull is usually significantly larger than the union of the FEMS. The result set of the database is only very occasionally smaller than that of the accelerator (when a very large number of records qualify for many different structure elements). Sorting in structure elements (which uses a significant amount of resources)

never takes place in the accelerator scenario. With the improvements in SAP Notes 1074953 and 1118425 (mentioned above), the compression of FEMS as optimization in the accelerator part (merge of duplicate structure elements into one structure element with information to which structure element it belongs to) and the OLAP part (being able to handle this information) has been implemented. It reduces the result set which has to be transferred from the accelerator to the BW system.

FEMS compression has no impact on the size of accelerator indexes. The FEMS compression feature is set to "off" as a default value. You can switch it on by adding the following line in section [OLAP] in the configuration file "TREXIndexServer.ini" (see SAP Note 1118425, mentioned above):

```
fems_compression = on
```

There used to be a bug with revision 42 as well as a problem with non-cumulative key figures. This problem is solved with SAP Note 1172709 (BI 7.0 SP19) – BIA: Non-cml. Key Figures Incorr if FEMS Compression Active.

SAP Note 1239230 (BI 7.0 SP20) – BIA: Assign Conflict When Using RSDRI to Access InfoCube is also related to FEMS compression. If there is an accelerator index for a non-cumulative InfoCube and RSDRI is used for access, this note has to be implemented.

For more information, refer to SAP Notes 1118425 and 1074953 (mentioned above). This intelligent compression schema eliminates redundant data in query results on the accelerator side and thus reduces data transfer times and cost.

Please also see BIA BLOG Series Part VIII: Improvements and Major Parameter Changes for BI Accelerator for additional details on other parameters.

<https://www.sdn.sap.com/irj/scn/weblogs?blog=/pub/wlg/11895>

9.2 Package-Wise Data Read

If using revision 49 or higher, consider activating "Package-Wise Data Read" to ensure all data can be read (without memory overflow). This setting does not make performance faster but puts the accelerator data to be read by BW in packages of configurable sizes, therefore streamlining the transfer of data. For detailed information, refer to SAP Notes 1157582 – BIA 7.00: New Feature "Package Wise Read, 1002839 – BIA 7.00: Error Due to Huge Result Set of Query and 1018798 – Reading Large Data Volumes from BI Accelerator.

In this scenario, the data volume read from the accelerator is too large. The problem is not the volume on the accelerator server, but the volume of data transferred from the accelerator server to the BW system after the query filter conditions have been applied and the aggregation has been carried out. There are two important structural differences between the accelerator server and the database:

- 1) The data is transferred using a Remote Function Call (RFC). This interface is not designed for transferring mass amounts of data.
- 2) The data is transferred in one package. It is not read in small packages (status of development before BI SPS15, BI SP17 and Accelerator revision 44), as it is for the database where activity is cursor based.

Solution

With BI SPS15 (BI SP17) and accelerator revision 44, SAP delivers the possibility to read data package-wise via the RFC from the accelerator server. With this, large data volumes can be read from the accelerator (with the performance of the RFCs). This capability will ensure that if the query runs without memory problems against the database, it will at least also run against the accelerator. When computing a query, BW requests query result sets from the accelerator. Without package-wise read, huge query result sets transferred via RFC could lead to memory problems and performance decreases.

The size of the query result sets sent back from the accelerator to BW can be customized now to suitable packages by the accelerator feature package-wise read. BW reads the accelerator results in packages of configurable size (= chunk_size) as long as the accelerator indicates to BW there is still data to be retrieved. The chunk_size basically limits the number of records multiplied by the number of key figures to be transferred from the accelerator server to BW in one package.

Limitation

The functionality will not eliminate all cases, especially those where queries dump due to memory problems in the BW application server. In this case, change the query definition accordingly e.g. by a split of the query into several queries with smaller result sets using more filter criteria. This also leads to more readable results for the user. SAP recommends using BW as a real OLAP tool, which means filtering and expanding the query as required. If unable to reformulate and change the query to reduce the data volume that is to be read, deactivate the accelerator server for this individual query. Call transaction SE16, table RSRREPDIR, and specify the technical name of the query as COMPID to select the relevant records. Change field NOHPA from " " to "X" or refer to SAP Note 1161525 – Do Not Use BIA for Queries.

Note the performance of such a query should generally not deteriorate significantly if you switch to the database. This is because the majority of the work is done processing in the OLAP processor and not the data selection.

Enablement

To enable package-wise read and in order to customize the package size, set the parameter "chunk_size" in the section [OLAP] of the TREX Index Server configuration file to the recommended value of "2,000,000". See SAP Note 1157582 discussed at the beginning of this section. The value "0" means that no packaging will be done which corresponds to the default setting. If the values of parameters "max_result_size_one_index" and "max_result_size_merge" according to SAP Note 1002839 (discussed earlier in this section) are increased too much the query might be successful on the accelerator server when using the package-wise read but may cause memory issues on the BW server.

SAP Note 1159305 – Wrong Data With BIA Package-Wise Read and Non-cumulatives is related to the package-wise read functionality. Wrong data can occur for non-cumulative data when package-wise data reading of BIA data is switched on. The issue is solved in this SAP Note.

Also refer to SAP Note 1002839 – BIA 7.00: Error Due to Huge Result Set of Query (BIA 7.00 Revision 48 and higher) which makes controlling the maximum result size simpler for the accelerator.

9.3 Miscellaneous

For other issues with large data volumes not addressed by the enabling of FEMS compression or Package-Wise Data Read, consider the following options:

- Reduce the accelerator query size for complex cases. For more information, refer to SAP Note 1085745 – Reduction of BIA Query Size in Complex Cases
- Enhance the interface between the data manager and OLAP processor. For more information, refer to SAP Note 1091714 – Dynamic Data Table During Reading of Data
- For non-cumulative key figures, use variable selection criteria to limit the data being queried and limit the size of the validity table. Be sure to compress InfoCube and re-index periodically. For more information, refer to SAP Note 1160520 – BIA and Non-cumulative Queries
- For queries that “dump” large data volumes and require no further drill down, consider using a DSO (with DB indexes) and not an InfoCube (with an Accelerator index). Unlike a query executed against an InfoCube, DSOs do not perform joins at query runtime. This is only an option that can be considered but each design is unique. Results will vary with different designs
- Limit the query to only those key figures and characteristics required for the query
- Set the "Use Selection of Structure Elements" property for queries with hidden key figures (can be set in RSRT)

10. OTHER CONSIDERATIONS

10.1 Aggregates

Prior to SAP NetWeaver BW Accelerator, aggregates provided the most effective way to improve query performance. Customers deploying the solution have removed most, if not all, aggregates. The reasons are obvious:

- Faster to roll-up data into the accelerator than aggregates.
- Change-runs require less time to complete if there are no (or very few) aggregates.
- Aggregates use up disk space.

The accelerator creates only one index for each InfoCube, so companies no longer need to waste administration resources maintaining multiple aggregates. Because the definition and creation of aggregates requires a detailed understanding of the business requirements, e.g., what queries are executed most frequently, what filter values are used, what drill-down path is used, etc., a high skill level is required. In addition, the ongoing task of maintaining them is labor intensive and the human resource costs for this approach to performance optimization are high. By contrast, both the initial indexing effort for an InfoCube and the subsequent maintenance effort for change runs and roll-ups are an order of magnitude lower. Furthermore, in most cases, the performance gain obtained is greater and more consistent with the accelerator.

The usage of aggregates needs to be monitored on a regular basis. While some aggregates may no longer be needed, new aggregates are required if there are any changes to the underlying cube, queries, or user behavior.

In short, aggregates have a negative impact on the storage requirements, data latency, and administration of BW. Please note that while it is fine to have both active aggregates and an accelerator index for a particular cube, the roll-ups and change runs are only faster if you have eliminated the aggregates.

Using the accelerator instead of aggregates has an impact on how cubes and queries can be designed. Before focusing on design issues, let's briefly discuss whether there are still reasons to keep (active) aggregates. Consider the following statements

1. If you already have a good response time, there is no need to delete the aggregates and create an accelerator index.
2. An aggregate may provide better performance than an accelerator index.

The first statement can easily be dismissed. If you do not have any aggregates and the performance is good, then you do not need to create an accelerator index. There are some customers that consider creating indexes even in this case, since the solution's flexibility can provide overall opportunities to lower TCO. For instance, with the accelerator the response time improvement is consistent. If the system is not under heavy load, response times for a given query are constant whenever it is executed. Also, similar queries are answered in similar times. By contrast, response times with aggregates often vary unpredictably, depending on whether an existing aggregate exactly matches a query or whether a database optimizer finds a perfect index for the query. The problem with relying on a database optimizer is that it is a "black box" product containing proprietary algorithms that impact performance in ways a user cannot anticipate. With the accelerator, users learn to expect fast responses to new queries. This encourages ad hoc reporting and creative analysis of the data. This in turn extracts the added value, not available with aggregates, hidden in the huge amounts of data generated by a modern business.

If aggregates exist, performance issues exist. Since the impact the accelerator has on roll-ups, change runs, disk space, maintenance, and support is significant, serious consideration needs to be given to delete many, if not all, aggregates and utilize the accelerator to its full potential

The second bullet discusses better performance with aggregates as opposed to the accelerator. This can be true but certainly is not the rule. It's alright to maintain an aggregate in these cases but understand the flexibility the accelerator provides which could potentially be forfeited if an aggregate is used. These are case by case decisions that need to be made. Is losing a few seconds in performance on a given cube or query worth sacrificing if aggregates do not have to be maintained?

Aggregates that should not be deleted are those that have been created in order to improve the performance of OpenHub. The accelerator was designed to improve the performance of ad hoc reporting - not data extraction from an InfoCube. In other words, the OpenHub interface will read the data from an InfoCube (or aggregate) even if an active index exists.

Aggregates are not entirely obsolete with the accelerator as it is possible to have active aggregates and an accelerator index for the same InfoCube. However, aggregates are typically not needed for a given InfoCube if indexed in the accelerator since queries bypass the database and the aggregates. Fewer aggregates will result in faster data load processes (aggregate change runs and roll-ups) and reduce the need for DASD. Most customers drop or deactivate their aggregates within weeks of going live with the accelerator. Consider keeping some aggregates after Go-Live.

- Until comfortable with the accelerator's stability, consider keeping aggregates for mission critical queries. If the accelerator is down, aggregates will be used (if applicable) and provide better query performance compared to reading data from the InfoCube/database.
- Aggregates are still needed for cubes not indexed with the accelerator.
- Some customers keep Basis aggregates (in the beginning). Basis aggregates do not contain navigational attributes and do not affect the change run.

10.2 OLAP Cache

SAP recommends keeping the OLAP cache switched on for InfoProviders with the accelerator indexes. The cache settings remain the same as for InfoProviders with aggregates. OLAP Cache can reduce the load from the accelerator and contribute to its scalability. When a query is executed with the accelerator it first checks to see if Information Broadcasting is used. Secondly, it checks the OLAP Cache and lastly the accelerator. If queries with OLAP-intensive operations are present, this would be a good use instead of the Accelerator

10.3 Other Processes

The accelerator removes workload from the database and minimizes the likelihood of other processes competing for the same database resources e.g. data loads. Note that running master or transaction data loads in parallel to the accelerator indexing jobs might create a performance issue. However, if process chains are set up properly, the accelerator indexing tasks (roll-ups) should always happen after completion of the actual load.

11. FUTURE OF SAP NETWEAVER BW ACCELERATOR - ROADMAP

The first release of the accelerator has been very successful and feedback has been extremely positive. Customers who became early adopters and have made the solution an integral part of their environment cannot imagine a scenario without it. However, like all solutions early in their lifecycle, numerous opportunities to grow the solution have been identified and will become part of the product in future releases.

With the next release, the accelerator will become more of an analytic engine and will be sold as two separate products. SAP NetWeaver BW Accelerator 7.2 (the follow-on to the product discussed in this document) will be the add-on to the SAP NetWeaver BW software and SAP BusinessObjects Accelerator will be targeted to those customers with SAP BusinessObjects data marts or WebI applications looking for greater acceleration.

For more detail on the roadmap of the accelerator solutions, please follow the link:

<https://www.sdn.sap.com/irj/sdn/go/portal/prtroot/docs/library/uuid/3604c604-0901-0010-f0aa-b37378495537>

11.1 SAP NetWeaver BW Accelerator 7.2

One of the biggest advancements with this release will be the ability to load data from a DSO to an InfoCube without database persistency, i.e. the fact data will be stored only on the accelerator. Also, the ability to load a DSO directly to this type of InfoCube with a DTP will exist. This will provide customers with greater flexibility in their warehouse by giving the option to bypass a data persisting InfoCube layer if desired, which will remove latency from the load window.

Just as important, is the shift to move more functionality off of the BW analytical engine and into the accelerator. With the next release, functionality such as top "n" and bottom "n" analysis, Multi-Provider aggregation and F4 Value Help will all be done in the accelerator, making the performance for these queries as fast as other queries which excel today with the current solution.

11.2 SAP BusinessObjects Accelerator

The SAP BusinessObjects Accelerator will be a stand alone appliance. It can be used in conjunction with any relational data sources with the help of Data Services. With this solution, reliable performance can be expected when querying on data from typical Infosets, MS SQL Server, SAP NetWeaver BW (native), SAP Business Suite (SAPI + native), Excel files, etc. With great query performance, additional agility and the lowering of TCO, this product is sure to make life for IT easier as the time needed to turn ideas from business users into real solutions should certainly decrease.

12. SUMMARY

To summarize, the SAP NetWeaver BW Accelerator value proposition is that deployment brings two significant benefits:

- **User productivity** is increased due to the improved query performance, more stable response times, high scalability and robustness, and the opportunities created by enhanced usability.
- **The total cost of operation** for analytic reporting is reduced by savings in aggregate maintenance, system consolidation, longer write-downs for existing hardware, ease of reporting, fewer resources required to support the system and simplified data modeling.

In addition to the increases in productivity as well as lowering the total cost of ownership, the accelerator will become an even more integral piece of the Business Intelligence environment in future releases of SAP NetWeaver BW as well as SAP BusinessObjects.

www.sdn.sap.com/irj/sdn/howtoguides