Redbooks

ibm.com/redbooks

Implementation Best Practices for IBM DB2 BLU Acceleration with SAP BW on IBM Power Systems

Dino Quintero

Yukiko Itaya

Speitim Velic

Adriana Melges Quintanilha Weingart

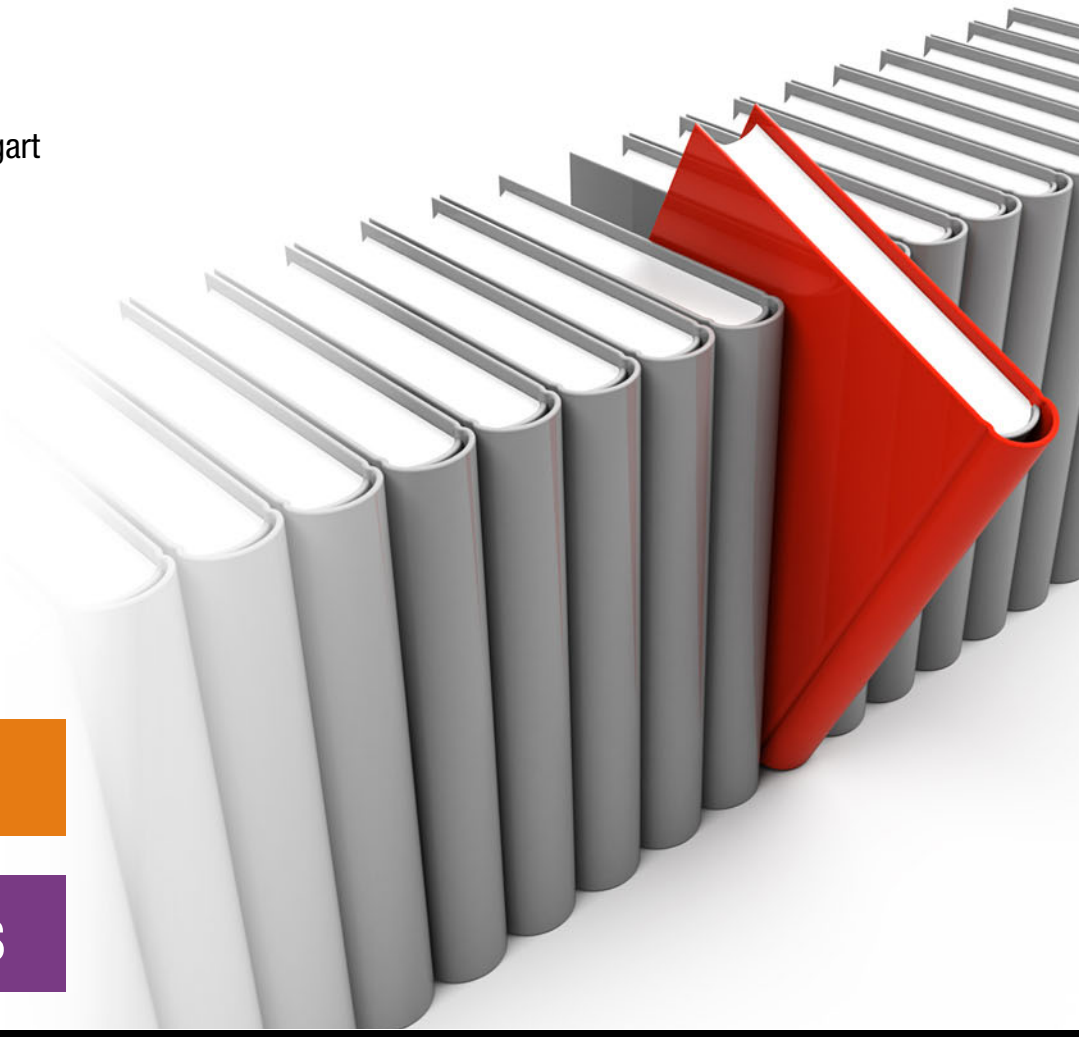

Analytics

Power Systems

IBM

Redbooks

International Technical Support Organization

Implementation Best Practices for IBM DB2 BLU Acceleration with SAP BW on IBM Power Systems

May 2015

SG24-8258-00

Note: Before using this information and the product it supports, read the information in “Notices” on page v.

First Edition (May 2015)

This edition applies to AIX 7.1 TL03, DB2 10.5 FP04.

© Copyright International Business Machines Corporation 2015. All rights reserved.
Note to U.S. Government Users Restricted Rights -- Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

Contents

Notices v
Trademarks vi

IBM Redbooks promotions vii

Preface ix
Authors ix
Now you can become a published author, too! x
Comments welcome xi
Stay connected to IBM Redbooks xi

Chapter 1. Introducing SAP NetWeaver Business Warehouse with IBM DB2 BLU 1
1.1 Current challenges with analytic systems 2
1.2 Customer needs 3
1.3 IBM DB2 with BLU Acceleration 4
1.3.1 Lower operating costs 6
1.3.2 Extreme performance 7
1.3.3 Hardware optimization 7
1.4 DB2 BLU proof of concept (PoC) results 10
1.5 Archiving with NLS 11
1.6 Multi-temperature storage group 13

Chapter 2. Requirements for an SAP NetWeaver BW environment 17
2.1 Introduction to SAP BW 18
2.2 DB2 with BLU Acceleration and SAP BWA 19
2.3 SAP Business Warehouse enhancements with DB2 10.5 Cancun Release (FP4) 19
2.4 System requirements and preparation 20
2.4.1 Software requirements 21
2.4.2 SAP DB2 10.5 bundle and SAP/IBM DB2 license 21
2.4.3 Restrictions for using BLU Acceleration 22
2.4.4 General sizing approach for DB2 BLU 22
2.4.5 Storage requirements 23
2.4.6 Platform requirements 24

Chapter 3. Optimizing IBM Power Systems for SAP BW using IBM DB2 BLU 25
3.1 TurboCore mode of the IBM POWER7 26
3.2 Start the most important LPAR for memory affinity 27
3.3 Set DB2_RESOURCE_POLICY for memory affinity 29
3.4 Dedicated logical partition for stable performance 30
3.5 Active memory expansion is not recommended in this case 34
3.6 AIX tunables 34
3.7 Sizing approach with IBM Power Systems 34

Chapter 4. SAP BW landscape scenario with IBM DB2 BLU Acceleration 37
4.1 Design SAP BW landscape 38
4.1.1 Hardware resources for SAP BW landscape with DB2 BLU 38
4.1.2 Considerations for the preproduction systems (DEV and QAS) 39
4.1.3 Consideration for the production system (PRD) 40
4.1.4 DB2 database layout for BLU 41

© Copyright IBM Corp. 2015. All rights reserved.

4.1.5 Migrate SAP BW to Power Systems with AIX 42
4.2 Implementation of SAP BW with DB2 BLU 42
4.2.1 Enabling adaptive row compression 42
4.2.2 Enabling automatic storage management during the installation 45
4.2.3 Enabling ASM after the SAP BW installation 46
4.2.4 Enabling intrapartition parallelism for BLU 51
4.2.5 DB2 Cancun parameter settings for BLU 52
4.2.6 SAP BW RSADMIN parameters for DB2 BLU Acceleration 53
4.2.7 Multi-temperature storage groups for DB2 temp table spaces 53
4.3 Administration and operation with SAP DB2 BLU 56
4.3.1 Transport system 56
4.3.2 Backup and restore 57
4.3.3 System copies 57

Chapter 5. Adding IBM DB2 Near-Line Storage to archive cold temperature SAP BW data 59
5.1 NLS scenario for SAP BW landscape 60
5.1.1 Design for SAP NLS 61
5.1.2 Minimum resources for SAP NLS 61
5.1.3 Considerations for the preproduction systems (NDE and NQA) 62
5.1.4 Consideration for the production system (NPR) 62
5.2 Implementing the Near-Line Storage feature 63
5.2.1 Installation of the NLS database NDE 63
5.2.2 Configure the new database connection in the DBACockpit 66
5.2.3 Configure the NLS connection in SPRO 67

Related publications 69
IBM Redbooks 69
Other publications 69
Online resources 69
Help from IBM 70

Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:
IBM Director of Licensing, IBM Corporation, North Castle Drive, Armonk, NY 10504-1785 U.S.A.

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law: INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM websites are provided for convenience only and do not in any manner serve as an endorsement of those websites. The materials at those websites are not part of the materials for this IBM product and use of those websites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs.

© Copyright IBM Corp. 2015. All rights reserved.

Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. These and other IBM trademarked terms are marked on their first occurrence in this information with the appropriate symbol (® or ™), indicating US registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of IBM trademarks is available on the Web at http://www.ibm.com/legal/copytrade.shtml

The following terms are trademarks of the International Business Machines Corporation in the United States, other countries, or both:

AIX®
DB2®
DS8000®
FlashSystem™
IBM®
IBM FlashSystem®
Micro-Partitioning®
POWER®
Power Architecture®
POWER Hypervisor™
Power Systems™
POWER6®
POWER7®
POWER7 Systems™
POWER7+™
POWER8™
PowerHA®
PowerVM®
pureScale®
Redbooks®
Redbooks (logo) ®
Storwize®
System p®
XIV®

The following terms are trademarks of other companies:

Intel, Intel logo, Intel Inside logo, and Intel Centrino logo are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

Linux is a trademark of Linus Torvalds in the United States, other countries, or both.

Windows, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Other company, product, or service names may be trademarks or service marks of others.

Find and read thousands of IBM Redbooks publications

- Search, bookmark, save and organize favorites
- Get up-to-the-minute Redbooks news and announcements
- Link to the latest Redbooks blogs and videos

Get the latest version of the Redbooks Mobile App

iOS **Download Now** Android

Promote your business in an IBM Redbooks publication

Place a Sponsorship Promotion in an IBM® Redbooks® publication, featuring your business or solution with a link to your web site.

Qualified IBM Business Partners may place a full page promotion in the most popular Redbooks publications. Imagine the power of being seen by users who download millions of Redbooks publications each year!

ibm.com/Redbooks
About Redbooks → Business Partner Programs

THIS PAGE INTENTIONALLY LEFT BLANK

Preface

BLU Acceleration is a new technology that has been developed by IBM® and integrated directly into the IBM DB2® engine. BLU Acceleration is a new storage engine along with integrated run time (directly into the core DB2 engine) to support the storage and analysis of column-organized tables. The BLU Acceleration processing is parallel to the regular, row-based table processing found in the DB2 engine. This is not a bolt-on technology nor is it a separate analytic engine that sits outside of DB2. Much like when IBM added XML data as a first class object within the database along with all the storage and processing enhancements that came with XML, now IBM has added column-organized tables directly into the storage and processing engine of DB2.

This IBM Redbooks® publication shows examples on an IBM Power Systems™ entry server as a starter configuration for small organizations, and build larger configurations with IBM Power Systems larger servers. This publication takes you through how to build a BLU Acceleration solution on IBM POWER® having SAP Landscape integrated to it.

This publication implements SAP NetWeaver Business Warehouse Systems as part of the scenario using another DB2 Feature called *Near-Line Storage* (NLS), on IBM POWER virtualization features to develop and document best recommendation scenarios.

This publication is targeted towards technical professionals (DBAs, data architects, consultants, technical support staff, and IT specialists) responsible for delivering cost-effective data management solutions to provide the best system configuration for their clients' data analytics on Power Systems.

Authors

This book was produced by a team of specialists from around the world working at the International Technical Support Organization (ITSO), Poughkeepsie Center.

Dino Quintero is a Complex Solutions Project Leader and an IBM Senior Certified IT Specialist with the ITSO in Poughkeepsie, NY. His areas of knowledge include enterprise continuous availability, enterprise systems management, system virtualization, technical computing, and clustering solutions. He is an Open Group Distinguished IT Specialist. Dino holds a Master of Computing Information Systems degree and a Bachelor of Science degree in Computer Science from Marist College.

Yukiko Itaya is an IT Specialist with IBM Japan Systems Engineering Co., Ltd. (ISE) under GTS in Japan, providing technical support on DB2 for Linux, UNIX, and Windows. She has 12 years of experience in technical support for IBM Power Systems, IBM AIX®, IBM PowerHA®, and DB2 for Linux, UNIX, and Windows and she is especially focusing on BLU Acceleration. She has worked with several major customers in Japan implementing DB2 for Linux, UNIX, and Windows and has conducted workshops for IBM staff in Japan.

© Copyright IBM Corp. 2015. All rights reserved.

Speitim Velic is an IT Specialist working in Germany for BWI Systeme GmbH (an IBM subsidiary). As a certified Technical Consultant for SAP NetWeaver and a Technology Consultant for OS/DB migration for SAP Systems, he has worked in a multitude of SAP areas since 2001. He holds a Master in Computer Science from 2003. His areas of expertise include also managing AIX Platform and DB2 Databases in an SAP Context.

Adriana Melges Quintanilha Weingart is a Certified Expert IT Specialist working in Brazil for IBM. She has 17 years of experience in SAP, with main areas of expertise in Performance and System Monitoring. Adriana has worked in several Brazilian and Global clients in Banking, Industry, Beverage, and Food Segments as an SAP Basis Specialist.

Thanks to the following people for their contributions to this project:

David Bennin, Richard Conway
IBM International Technical Support Organization, Poughkeepsie Center

Scott Maberry, Augie Mena, Bret Olszewski
IBM US

Douglas Gibbs, Neil Graham, Haider Rizvi, Charlie Wang
IBM Canada

Olaf Depper, Karl Fleckenstein, Mirco Malessa
IBM Germany

Uwe Bellinghausen, Hans-Peter Gronimus, Jochen Kellner, Dirk Leonhardt, Kerstin Reichelt, Irmgard Schröter
BWI Systeme GmbH (an IBM Subsidiary)

Mayumi Hirano, Oyama Junichi, Tetsuya Shirai, Kurokawa Yoshiaki
IBM Japan

Vinícius Cosmo Cardoso, Cleiton Soares Freire, Soraya Alves Schlossarecke Mallavazzi, Marco Poma
IBM Brazil

Now you can become a published author, too!

Here's an opportunity to spotlight your skills, grow your career, and become a published author—all at the same time! Join an ITSO residency project and help write a book in your area of expertise, while honing your experience using leading-edge technologies. Your efforts will help to increase product acceptance and customer satisfaction, as you expand your network of technical contacts and relationships. Residencies run from two to six weeks in length, and you can participate either in person or as a remote resident working from your home base.

Find out more about the residency program, browse the residency index, and apply online at:

ibm.com/redbooks/residencies.html

Comments welcome

Your comments are important to us!

We want our books to be as helpful as possible. Send us your comments about this book or other IBM Redbooks publications in one of the following ways:

- Use the online **Contact us** review Redbooks form found at:

  ibm.com/redbooks

- Send your comments in an email to:

  redbooks@us.ibm.com

- Mail your comments to:

  IBM Corporation, International Technical Support Organization
  Dept. HYTD Mail Station P099
  2455 South Road
  Poughkeepsie, NY 12601-5400

Stay connected to IBM Redbooks

- Find us on Facebook:

  http://www.facebook.com/IBMRedbooks

- Follow us on Twitter:

  https://twitter.com/ibmredbooks

- Look for us on LinkedIn:

  http://www.linkedin.com/groups?home=&gid=2130806

- Explore new Redbooks publications, residencies, and workshops with the IBM Redbooks weekly newsletter:

  https://www.redbooks.ibm.com/Redbooks.nsf/subscribe?OpenForm

- Stay current on recent Redbooks publications with RSS Feeds:

  http://www.redbooks.ibm.com/rss.html

1

Introducing SAP NetWeaver Business Warehouse with IBM DB2 BLU

This chapter introduces you to the most common business warehousing scenarios, and the challenges that such scenarios bring to the technology infrastructure supporting them. Some of today's challenges involve how to create solutions that get better, are faster, and more stable, such that there is a process to guarantee the information provided to the user.

In addition to providing an overview of the solution, this chapter also describes the approach that is taken in this IBM Redbooks publication, and the reasons that led to the selection of the following solutions: *SAP NetWeaver Business Warehouse* (known as SAP NetWeaver BW, or SAP BW) running on IBM DB2 for Linux, UNIX, and Windows (DB2) with BLU Acceleration and AIX, with Near-Line Storage (NLS) and multi-temperature storage groups.

This publication's purpose is not to define the final nor the best solution for the client's needs, but describes a solution and the benefits to help the clients provide the best infrastructure for daily operation of the client's system.

The following sections are described in this chapter:

- Current challenges with analytic systems
- Customer needs
- IBM DB2 with BLU Acceleration
- DB2 BLU proof of concept (PoC) results
- Archiving with NLS
- Multi-temperature storage group

© Copyright IBM Corp. 2015. All rights reserved.

1.1 Current challenges with analytic systems

Checking across several environments and clients using SAP solutions as a platform for Business Intelligence, it is possible to find different infrastructure scenarios with different combinations of hardware and software, and with many features available.

In the scenarios, it is possible to find some common points such as:

1. Data increasing: As a result of the world's actual need for information (saved and ready for usage) systems are growing very fast. Business growth adds to these complex solutions in addition to the new possibilities allowed by the solution and the new requirements from the day-to-day business operations.
2. Quickest ways to access the data: The speed required in today's IT environments is a critical challenge, considering the performance in a query or when managing the (large amounts) data or even when managing a system, and sometimes with concurrent tasks. For example, in a business intelligence (BI) system, there is a point to consider: While the ETL is running, it is necessary to extract, in parallel, reports and there is no time to lose waiting for them.
3. Costs increase:
   a. As a result of the data growth and the need for quick system responses, the investments to make more hardware resources available is a constant concern for the clients, in additional to the already ongoing maintenance costs.
   b. Also, software investments are a source of concern as the software is getting back level, in addition to the ongoing maintenance costs, and the need for new software features. At the same time, it is necessary to keep the production environment working, and not only working, but providing good response time (quality of service) to the business workloads.

Figure 1-1 shows you the IT costs distribution according to *IDC July 2010: The Business Value of Large-Scale Server Consolidation:*

http://www-05.ibm.com/innovation/cz/leadership/pdf/POL03073USEN.pdf

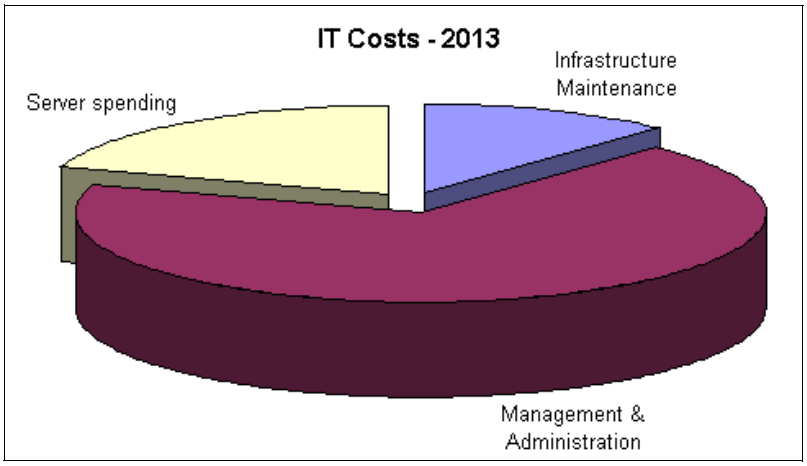

Figure 1-1 IT costs distribution

1.2 Customer needs

Considering today's challenges and scenarios, it is not difficult to decipher what the clients are looking for the most: Faster system response times, cost reduction, and new technology features. Figure 1-2 presents some of the features that make up the clients' wish lists.

	Performance	Cost Reduction	Optimization
Simplicity		X	
Data Compression	X	X	X
Deep Hardware Instruction Exploitation	X		X
Parallelism	X		X
Memory Caching	X		X

Figure 1-2 Clients' wish list

Some of the main reasons behind the clients' wish list are explained as follows:

1. Reducing costs is an important item to consider for all IT environments. Usually IT is not the major concern the clients have for their business. However, the possibility of reducing costs allows IT to get better, and improve without impacting the business. The following actions can reduce IT costs, and can have direct impact on the business:
   - Simplicity: Keep it simple. Sometimes, there is no way but to invest money on new solutions or development to attend to some business requirements. However, whenever possible, reuse the solutions, or better to use the actual solutions, saving time (and money) to install, configure, and later, maintain and administer. Reusing the actual solution and existing tools is a good way to keep things simple.
   - Data compression: A great part of IT costs is related to disks and how to store data. Besides the traditional archiving and housekeeping solutions, any other possibility that allows saving space is welcomed as a way to save costs. Data compression can work to save disk space, besides the power, cooling, and management of the storage units.
2. In a world where information and speed are mandatory, IT systems' performance is a critical point for business. Several factors can be considered for improving system capacity and allowing faster response times:
   - Data compression techniques used to compress data before storing it, or to decrease the amount of data transmitted within the enterprise.
   - Better hardware usage, such as deep hardware instruction exploitation, parallelism, and memory caching allow existing technologies to work faster and more efficiently.

In other words, extreme performance, lower operating costs, and hardware optimization can summarize the client's needs. All other needs are usually related to these, and are shown in Figure 1-3.

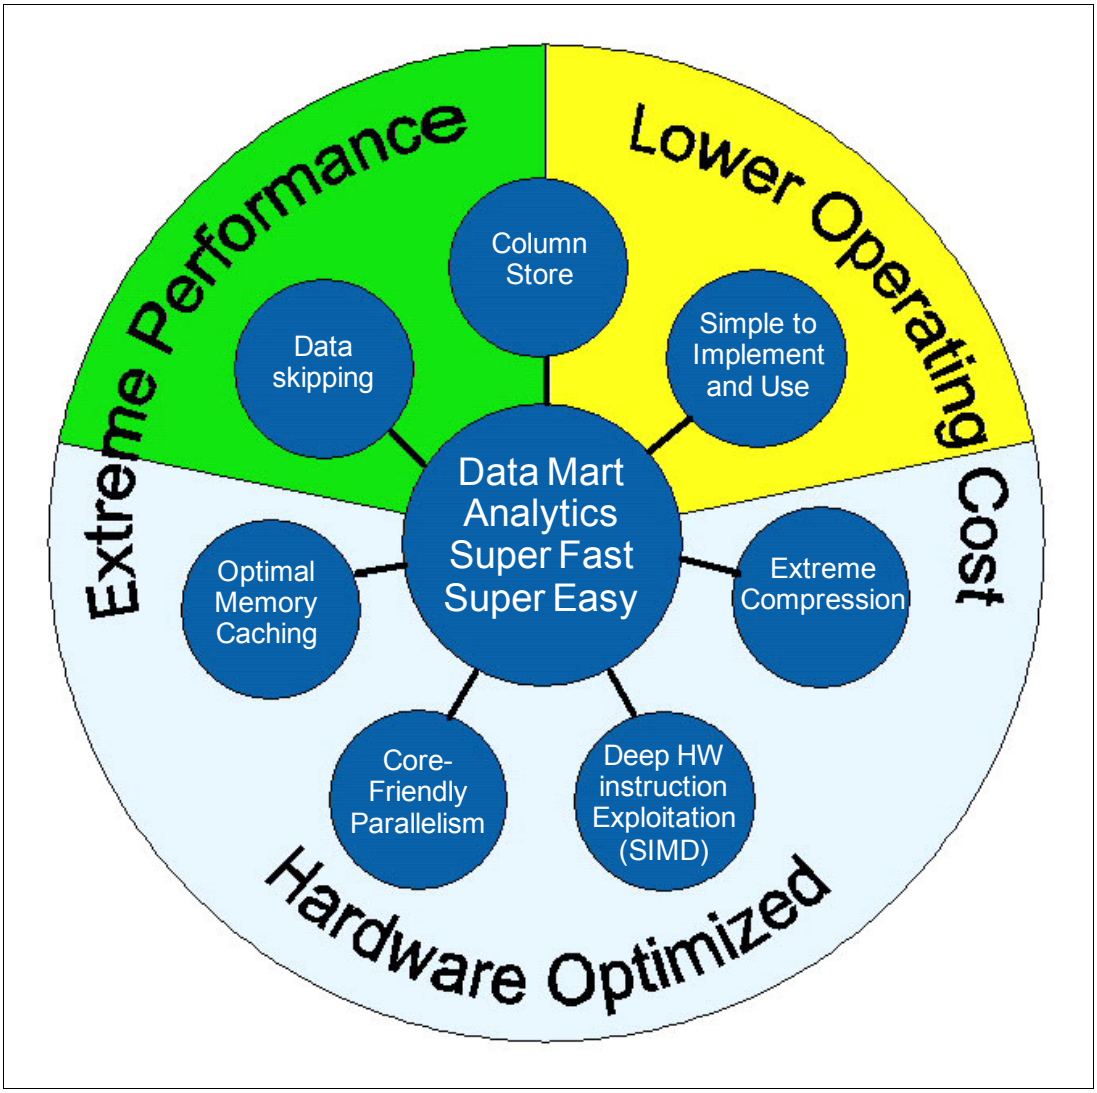

Figure 1-3 What are clients looking for?

1.3 IBM DB2 with BLU Acceleration

BLU Acceleration is one of the most significant advances in technology that has ever been delivered in DB2 and arguably in the database market in general. Available with the IBM DB2 10.5 release, BLU Acceleration delivers unparalleled performance improvements for analytic applications and reporting by using dynamic in-memory optimized columnar technologies.

Although BLU Acceleration is important new technology in DB2, BLU Acceleration is built directly into the DB2 kernel. BLU Acceleration is not only an extra feature; it is a part of DB2, and every component of DB2 is aware of BLU Acceleration. BLU Acceleration still uses the same storage unit of pages, the same buffer pools, and the same backup and recovery mechanisms.

Note: The term “BLU” does not stand for anything in particular. It may be related to the Blink Ultra project, which resulted on the BLU Acceleration, or even be an indirect play with the IBM traditional corporate nickname “Big Blue”.

Figure 1-4 presents a general overview of DB2 with BLU Acceleration, which is discussed more in the following sections.

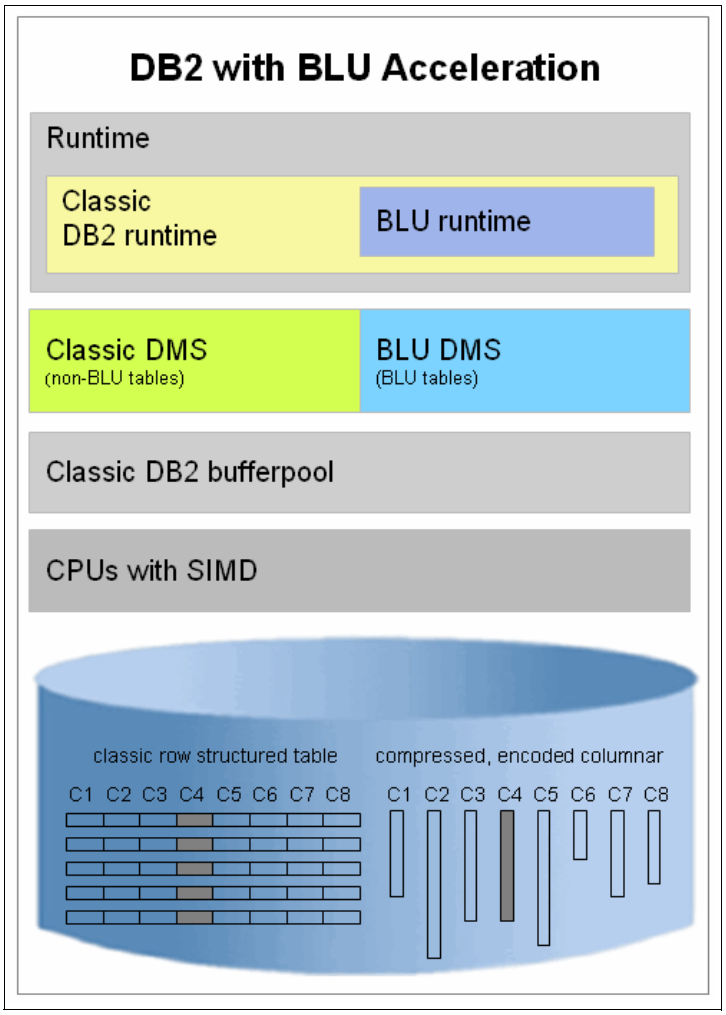

Figure 1-4 DB2 with BLU Acceleration: General overview

For more information about BLU Acceleration, see *Architecting and Developing DB2 with BLU Acceleration*, SG24-8212:

http://www.redbooks.ibm.com/abstracts/sg248212.html

Typical experiences using DB2 with BLU Acceleration shows good approaches and results such as:

- Simple to implement and use.
- Helps performance gains of about 10x - 20x.
- Helps storage savings versus decompressed data of about 5x - 20x.

Performance and response time of IT systems, specially BI (Business Information) systems when running reports are always a source of concern, and no matter what is done, these systems' performance and response times can always be improved.

1.3.1 Lower operating costs

IT costs are usually a source of discomfort in companies. IT is, with few exceptions, viewed as a cost and not as an investment. Considering this assumption, all the savings IT can provide with the maximum results are always welcomed.

Simple to implement and use

Keep it simple is a strong, almost mandatory concept in BLU Acceleration. Since its implementation, maintenance, and daily usability, the following keep-it-simple characteristics are present:

- One setting is necessary to optimize the DB2 system for BLU Acceleration. It is only necessary to set one database variable, when the database is used for analytic workloads, and optimize for optimal analytics performance.
- No additional workload and maintenance: No indexes, multi-dimensional clustering (MDC), statistics views, manual reorganization, or RUNSTATS (these last two are automated).
- All the features built into the DB2 kernel: SQL, language interfaces, administration, reusing DB2 process model, storage concepts, and utilities.
- Simple table creation and conversion.

Note: DB2 BLU Acceleration activation in SAP NetWeaver Business Warehouse (SAP BW) Systems is different and requires SAP procedures and recommendations to be followed. Chapter 2, "Requirements for an SAP NetWeaver BW environment" on page 17, Chapter 3, "Optimizing IBM Power Systems for SAP BW using IBM DB2 BLU" on page 25, Chapter 4, "SAP BW landscape scenario with IBM DB2 BLU Acceleration" on page 37, and Chapter 5, "Adding IBM DB2 Near-Line Storage to archive cold temperature SAP BW data" on page 59 illustrate how to integrate the DB2 BLU Acceleration feature.

Column store

The most basic and prominent feature of BLU Acceleration is the column-organized table type.

Column-organized tables stores each column on a separate set of pages on disk, reducing I/O needed for processing queries that are loaded into memory from disk. Following are the features that in the tests helped generate savings:

- Minimal I/O: By only performing I/O in the columns and values that match the query, and reducing the working set of pages during the query progress.
- Work performed directly in columns: By working on individual columns for predicate evaluations, joins, scans, and so on, and not materializing rows until necessary to build the result set.
- Improved memory density and extreme compression: By keeping columnar data compressed in memory and packing more data values into a small amount of memory or disk.
- Cache efficiency: By packing data into CPU cache-friendly structures.

By being able to store both row and column-organized tables in the same database, as illustrated in Figure 1-5, users can implement BLU Acceleration even in database environments where mixed online transaction processing (OLTP) and online analytical processing (OLAP) workloads are required. Again, BLU Acceleration is built into the DB2 engine. The SQL, optimizer, utilities, and other components are fully aware of both row-organized and column-organized tables at the same time.

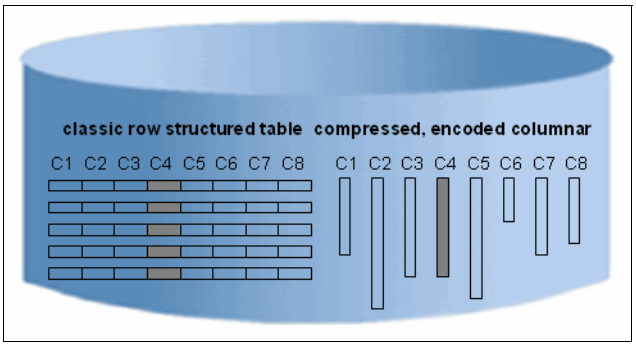

Figure 1-5 Classic row-structured table X compressed, encoded columnar table

1.3.2 Extreme performance

Besides lowering costs, providing performance for business operation is expected from the IT infrastructure. Columnar storage already helps this requirement by improving the I/O, memory, and cache efficiently, but the BLU Acceleration has another feature that allows this feature to be more significant: data skipping.

Data skipping

Data skipping avoids the unnecessary processing of irrelevant data, further reducing the I/O that is required to complete a query.

Automatic detection of large sections of data that does not qualify for a query, can be ignored. Data skipping is used for data in memory (buffer pool) as well as on disk and helps to significantly reduce I/O, memory, and CPU consumption.

This is how BLU Acceleration does data skipping: As data is loaded into column-organized tables, BLU Acceleration tracks the minimum and maximum values on ranges of rows in metadata objects that are called the *synopsis tables*. These metadata objects (or synopsis tables) are dynamically managed and updated by the DB2 engine without intervention from the DBA. When a query is run, BLU Acceleration looks up the synopsis tables for ranges of data that contain the value that matches the query. It effectively avoids the blocks of data values that do not satisfy the query, and skips straight to the portions of data that matches the query. The net effect is that only necessary data is read or loaded into system memory, which in turn provides dramatic increase in speed of query execution because much of the unnecessary scanning is avoided.

1.3.3 Hardware optimization

Lowering costs and providing optimal performance works together with hardware optimization. The most effective way the hardware is optimized and used to its full potential, it is possible to improve the performance and lower the costs for managing new hardware configurations.

BLU Acceleration provides the following features.

Extreme or adaptive compression

The column data is compressed with actionable compression, which preserves order so that the data can be used without decompression, resulting in storage and CPU savings and a significantly higher density of useful data that is held in memory.

This is possible because of the following features:

- Massive compression with approximate Huffman encoding, considering that the more frequent the value, the fewer bits it takes.
- Encoded values packed into bits matching the register width of the CPU.
- Late materialization, which is the ability to operate on the data while it is still compressed. Predicates and joins work directly on encoded values (actionable compression).

In addition to column-level compression, BLU Acceleration also uses page-level compression when appropriate. This helps further to compress the data based on the local clustering of values on individual data pages.

Because BLU Acceleration is able to handle query predicates without decoding the values, more data can be packed in the processor cache and buffer pools. These results in less disk I/O, better use of memory, and more effective use of the processor. Therefore, query performance is better and storage utilization is also reduced.

Deep hardware exploitation (SIMD)

BLU Acceleration optimizes the entire access to the hardware and its usage, seeking every opportunity (memory, CPU, and I/O). BLU Acceleration is designed to fully exploit all the computing resources provisioned to the DB2 server. This is possible by using *Single Instruction Multiple Data (SIMD)* capable CPUs.

SIMD instructions are low-level specific CPU instructions. DB2 can use a single SIMD instruction to get results from multiple data elements (perform equality predicate processing, for example) if they are in the same register.

DB2 has deep processor and memory exploitation, including AIX Workload Management (WLM) characteristics within its own workload policies.

Considering the processor exploitation, DB2 has:

- Deep exploitation of simultaneous multithreading (SMT).
- Key POWER value proposition with the ability to dispatch many threads.
- Decimal arithmetic that is performed directly on the DECFLOAT accelerator.

DB2 with BLU Acceleration has special algorithms that automatically take advantage of the built-in parallelism in the processors, if SIMD-enabled hardware is available. This is another feature in BLU Acceleration that allows using special hardware instructions that work on multiple data elements with a single instruction as shown in Figure 1-6 on page 9.

DB2 presents for memory exploitation the following characteristics:

- Autonomic detection and exploitation of POWER features such as larger pages and storage keys.
- Alternative page cleaning algorithms built into DB2 specifically for AIX performance boost.

DB2 has for storage exploitation the following characteristics:

- Uses asynchronous I/O and scatter/gather I/O, AIX concurrent I/O (CIO), and AIX direct I/O (DIO) interfaces.
- End-to-end I/O priorities (WLM, prefetcher, and AIX I/O prioritization).

More information about hardware exploitation is available at the following website:

http://www-01.ibm.com/software/data/db2/power-systems

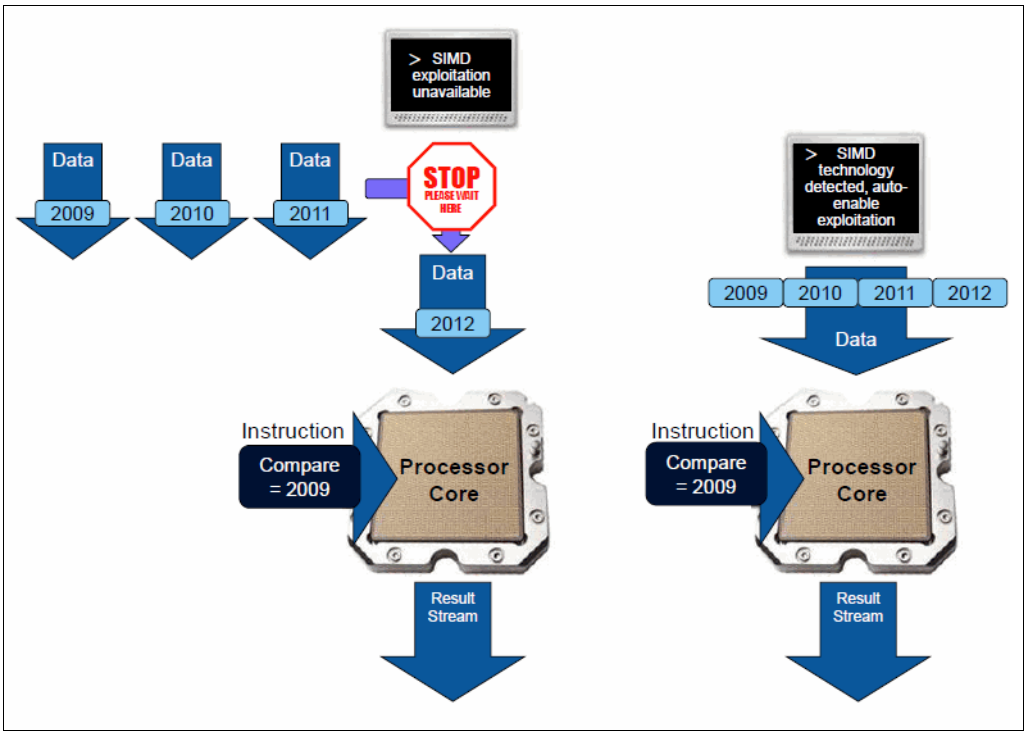

Figure 1-6 Use of SIMD processor with BLU Acceleration

Besides hardware optimization, BLU Acceleration contributes to lower the costs as it does not require users to acquire new hardware. In other terms, if a SIMD-enabled processor is not available in the current hardware, BLU Acceleration still functions and speeds query processing. In some cases, BLU Acceleration even emulates SIMD parallelism with its advanced algorithms to deliver similar SIMD results.

Core-friendly parallelism

BLU Acceleration is a dynamic in-memory technology. It makes efficient use of the number of processor cores in the current system, allowing queries to be processed by using multi-core parallelism and scale across processor sockets. The idea is to maximize the processing from processor caches, and minimize the latencies having to read from memory and, last, from disk.

Core-friendly parallelism is comprehensive algorithms that are designed to carefully place and align data that is likely to be revisited into the processor cache lines. This maximizes the hit rate to the processor cache, thus increasing effectiveness of cache lines.

Parallel vector processing with multi-core parallelism, single instruction, and multiple data parallelism provides improved performance and better utilization of available CPU resources.

Following are some physical attributes of the server:

- Queries on BLU tables are automatically parallelized.
- Fully uses the power of multiple CPU cores.
- Maximizes CPU cache efficiency to optimize cache lines.

Optimal memory caching

DB2 automatically adapts the way that it operates based on the organization of the table (row-organized or column-organized) being accessed.

BLU Acceleration uses some attributes to optimize the memory cache:

- New algorithms cache effectively in RAM (buffer pool).
- Data can be larger than RAM. There is no need to ensure that all data fits in memory.
- Separates caching algorithms for BLU/OLTP data:
  - BLU: Scan-friendly caching that minimizes I/O.
  - OLTP: LRU-based page cleaning that reclaims buffer pool space without regard to future I/O.

BLU Acceleration includes a set of big data-aware algorithms for cleaning out memory pools that are more advanced than the typical least recently used (LRU) algorithms that are associated with traditional row-organized database technologies. These BLU Acceleration algorithms are designed from the ground up to detect data patterns that are likely to be revisited, and to hold those pages in the buffer pool as long as possible. These algorithms work with DB2 traditional row-based algorithms.

1.4 DB2 BLU proof of concept (PoC) results

IBM DB2 10.5 delivers built-in a very powerful technology, the BLU Acceleration, which is based on column-organized tables instead of pages containing row data that presents:

- Adaptive compression and data skipping: By using row and column-organized tables together, getting the best of both worlds, BLU Acceleration helps deliver over 10x compression capabilities. This reduces power and cooling costs for storage, and also savings on backups, either in storage and in time required for the backups. Also, the I/O throughput is positively impacted and can provide performance savings.
- Hardware exploitation: BLU Acceleration also delivers 10 - 20x performance when using parallel vector processing, core-friendly parallelism, and scan-friendly memory caching. These features allow better hardware usage, and even savings or possibilities for growth without new costs.
- Dynamic in-memory columnar technology, in other words, column-organized tables that benefits analytic workloads: Using this feature, BLU Acceleration does not require auxiliary performance structures and procedures, such as indexes, materialized views (MQTs), multi-dimensional clustering (MDS), statistics gathering, or reorganization (which is done automatically to reclaim space without intervention by DBAs) to obtain the performance and response time reduction. These characteristics allow, besides the hardware (storage) savings, the management and administration efforts and costs to be reduced as well.
- Easy implementation and management: Still considering the built-in features of DB2, no new hardware or tool is necessary to be installed in order to activate the Acceleration. And the no need for auxiliary performance activities and less storage allows management savings on resources and time.

1.5 Archiving with NLS

IBM DB2 Near-Line Storage (NLS) is a solution for SAP BW systems. Its design allows online performance-enhancing while accommodating data growth, which are two important needs for systems that require reducing costs and retaining fast access to information.

With BLU Acceleration's in-memory columnar processing, it is possible to enable frequently-used data to be shifted from the SAP BW to DB2 NLS, allowing improvements in the entire system performance.

Using the near-line environment (or *near-online*, as frequently called), the NLS solution creates a middle ground between online and traditional archival systems.

Infrequently used data can be moved out of the online database to a separate database, reducing the size of the online database and consequently boosting its performance. Meanwhile, the data is kept in a DB2 database that enables organizations to capitalize on DB2 management capabilities. This means that the users can access information in the DB2 database more quickly than if the data was archived in another file system or on tape. The general implementation of the Near-Line Storage solution can be viewed in Figure 1-7 on page 12.

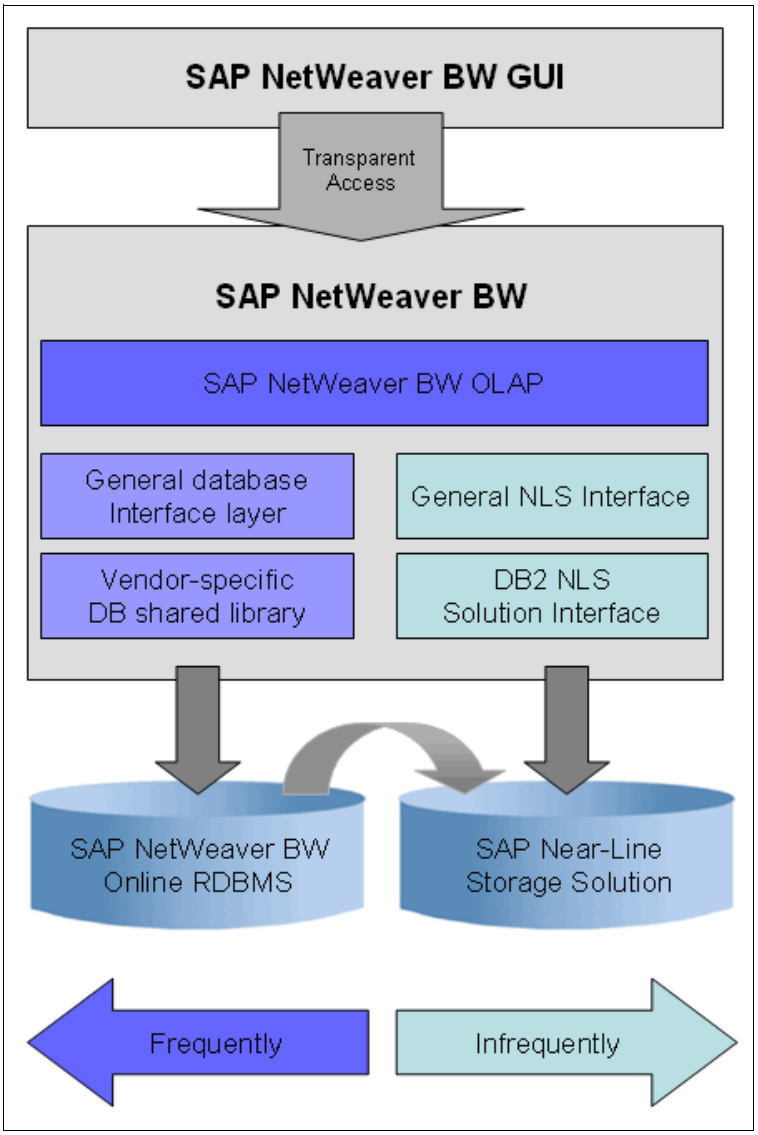

Figure 1-7 Basic setup for the Near-Line Storage solution

To implement these features, some assumptions need to be considered while using NLS:

- NLS saves infrequently accessed data.
- Slower query access on NLS data is acceptable.
- There are lower high availability requirements on NLS data.
- NLS data can be stored on less expensive hardware.

The NLS solution for SAP BW is integrated into the SAP data archiving strategy (SAP ADK). An SAP BW InfoProvider can be archived with a Data Archiving Process (DAP), which is extended for NLS.

According to the data importance and age, the data can be stored in different storage: online database, NLS, or classical archive, as shown in Figure 1-8.

	Online Database	Near-Line Storage (NLS)	Classic Archive
Frequently accessed data (actual)	✓		
Infrequently read data	⅄	✓	
Very rarely read data (aged)	✗	⅄	✓

✓ Optimal
⅄ Tolerable
✗ Unacceptable

Figure 1-8 Information lifecycle according to data importance and age

Chapter 5, “Adding IBM DB2 Near-Line Storage to archive cold temperature SAP BW data” on page 59 discusses and implements NLS.

1.6 Multi-temperature storage group

In a live system, all data is not always accessed in the same way or frequency. Some data is more frequently accessed than others. Even among data that is more accessed, there is data that is constantly accessed and data that is infrequently accessed. This data can be classified in the following ways:

- Hot data: Part of the data that is accessed often.
- Warm data: Another part of the data, frequently accessed, but not so much as the hot data.
- Cold data: The largest part of the data, which is accessed infrequently.
- Dormant data: Usually regulatory type or archival data that is accessed infrequently and never updated.

This classification is called *temperature* as shown in Figure 1-9, and refers to the fact that the data is not accessed the same way in a system.

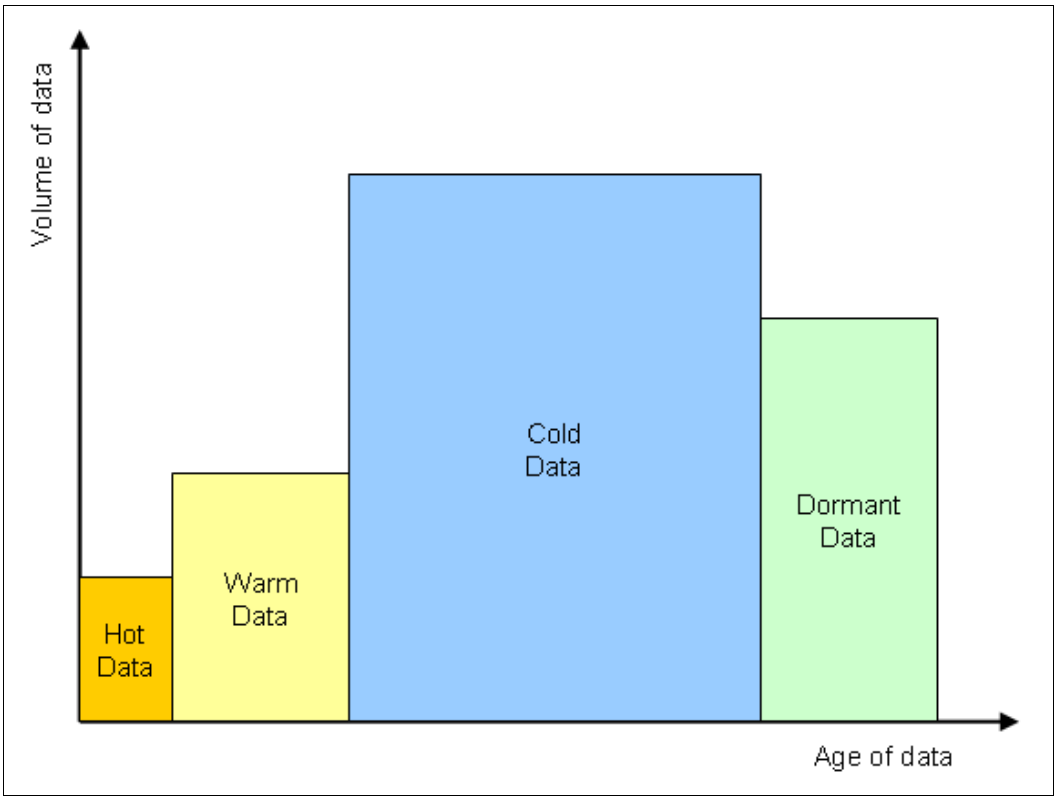

Figure 1-9 Data distribution across the temperature tiers

Usually, the multitier classification (hot, warm, cold, and dormant data) is not distinguished. A data management strategy that uses different storage media for data of different temperature is also called a *multi-temperature data management strategy*.

Considering large systems, this multi-temperature data management strategy is becoming even more required. DB2 for Linux, UNIX, and Windows 10.1 introduces a new feature, *the storage groups*, which allows this three-tier strategy to be implemented.

Storage groups are specific storage types, reflecting their characteristics in terms of performance and property of every automatic storage table space. In the simplest term, this is an extra layer between the table spaces and the automatic storage paths in the database as shown in Figure 1-10.

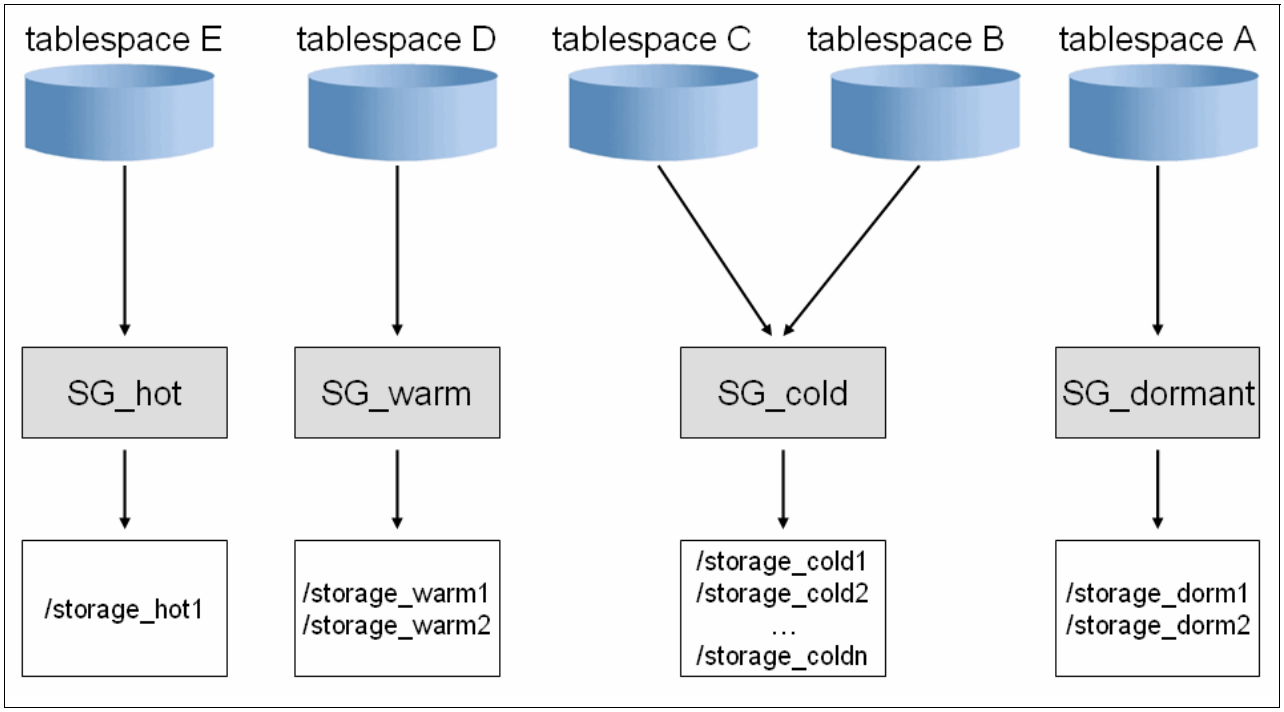

Figure 1-10 Storage groups

Multi-temperature storage groups are covered in more details in Chapter 4, “SAP BW landscape scenario with IBM DB2 BLU Acceleration” on page 37.

2

Requirements for an SAP NetWeaver BW environment

This chapter describes the requirements and recommendations for implementing SAP NetWeaver Business Warehouse (known as SAP NetWeaver BW, or SAP BW) using IBM DB2 with BLU Acceleration and Near-Line Storage (NLS) on IBM Power Systems.

This chapter discusses the following topics:

- Introduction to SAP BW
- DB2 with BLU Acceleration and SAP BWA
- SAP Business Warehouse enhancements with DB2 10.5 Cancun Release (FP4)
- System requirements and preparation

© Copyright IBM Corp. 2015. All rights reserved.

2.1 Introduction to SAP BW

SAP BW is a data warehouse system part of the SAP NetWeaver platform (provided by SAP) that integrates, transforms, and consolidates relevant business information from several sources, to provide an infrastructure that supports data evaluation and analysis.

Other SAP products such as the SAP Strategic Enterprise Management (SAP SEM), and SAP Business Planning and Consolidation (SAP BPC) are technologically based on SAP BW.

Some key capabilities of SAP BW are as follows:

- Buildup and management of the enterprise.
- Departmental data warehouses for business planning.
- Enterprise reporting and analysis.

SAP BW can extract master data and transactional data from various data sources and store them in the data warehouse. The SAP BW information model and data flow can be viewed in Figure 2-1.

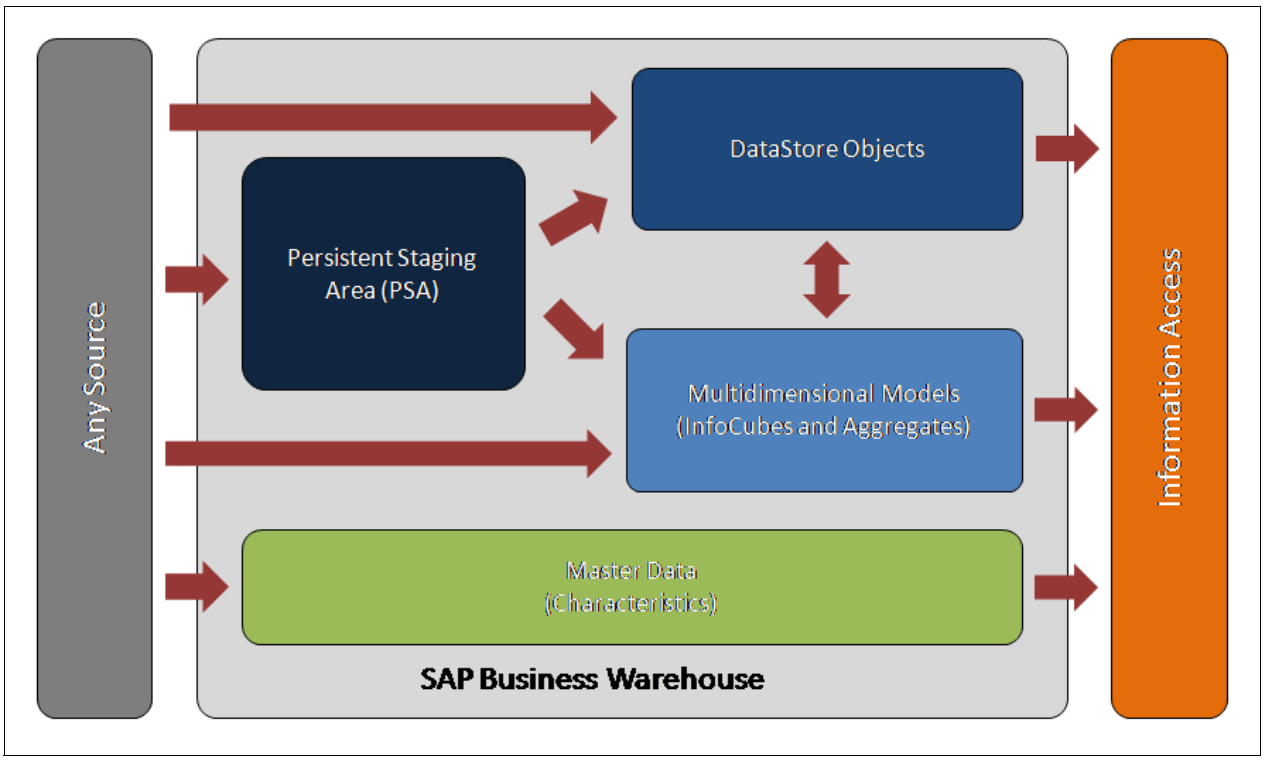

Figure 2-1 SAP Business Warehouse information model and data flow

The SAP NetWeaver central business warehouse platform for SAP customer functions are:

- Data comes into the system via extract, transform, and load (ETL).
- Users run BW reports on InfoProviders like InfoCubes and Operational Data Store.

For more information about SAP BW capabilities, see *Architecting and Developing DB2 with BLU Acceleration*, SG24-8212:

http://www.redbooks.ibm.com/abstracts/sg248212.html

2.2 DB2 with BLU Acceleration and SAP BWA

One possibility for the SAP BW Systems Acceleration is found in SAP with the SAP BW Accelerator (BWA), which is an approach for boosting the system performance based on SAP's search and classification engine (TREX) on specially configured hardware.

SAP BWA uses a TREX aggregation engine to improve the read performance of the queries on an InfoCube. SAP BWA is suited to situations in which relational aggregates or other means of improving performance (such as database indexes) do not suffice, or are too laborious.

In parallel, other options exist, such as the DB2 BLU Acceleration which SAP has certified SAP BW's use with IBM DB2 10.5's BLU Acceleration. The in-memory database for reporting and business intelligence allows the DB2 powerful in-memory optimized parallel vector processing columnar engine to accelerate their workloads without the need of extra hardware or a hardware upgrade.

DB2 BLU Acceleration provides the same results as SAP BWA, but uses less hardware resources, and simplifies administration due to the following factors:

- No separate appliance required with BLU (SAP BWA is a separate box).
- SAP BWA needs more administration.
- SAP BWA must be updated constantly with data from SAP BW.
- Data load requires more time (ETL) and causes extra workload on SAP BW.

Note: Users that have SAP ASL licenses (also known as OEM licenses) for DB2 can get BLU at no additional charge. Also, users who have Advanced Enterprise Server (AESE) (and who have support) can upgrade to 10.5 AESE and use BLU under SAP, and also get BLU Acceleration at no extra charge. For more details about SAP licensing, see 2.4.2, "SAP DB2 10.5 bundle and SAP/IBM DB2 license" on page 21.

2.3 SAP Business Warehouse enhancements with DB2 10.5 Cancun Release (FP4)

The use of DB2 BLU Acceleration is supported for SAP BW and SAP SEM as stated by the *SAP Notes 1819734* and *1825340* as shown in Table 2-1.

Table 2-1 SAP notes supporting SAP for BLU Acceleration

SAP note	Description
1819734	DB6: Use of BLU Acceleration. Describes prerequisites and restrictions for using BLU Acceleration in the SAP environment (includes corrections).
1825340	DB6: Use of BLU Acceleration with SAP BW and applications based on SAP BW. Describes prerequisites and the procedure for enabling BLU Acceleration in SAP BW.

Since SAP BW 7.0, all versions of SAP BW have integrated the features present in DB2 BLU Acceleration depending on the fix pack level DB2 has as explained in Figure 2-2.

	DB2 10.5 FP1	DB2 10.5 FP3aSAP	DB2 10.5 FP4
Support for DB2 BLU Acceleration for SAP BW	SAP BW - Standard InfoCube - Non-Cumulative InfoCube - InfoCubes in Multi-Providers - Semantical partitioned InfoProviders Near-Line Storage	- Basic cumulative InfoCubes - Basic non-cumulative InfoCubes - MultiProvider with underlying column-organized basic InfoCubes - Semantically partitioned InfoCubes (SAP BW 7.3 and higher only) NOTE: Not possible to create column-organized real-time InfoCubes or transactional InfoCubes.	- Standard DataStore objects - Write-optimized DataStore objects - InfoObjects (characteristics, time characteristics, and units) - Persistent staging areas (PSAs) - Flat InfoCubes (SAP BW 7.4 SP08 and higher only)

Figure 2-2 DB2 with BLU Acceleration support for SAP BW Systems

2.4 System requirements and preparation

In order to activate and use the DB2 BLU Acceleration features in an SAP BW System (or SAP SEM System), it is necessary to follow some SAP recommendations and requirements. *SAP Note 1819734*[1] describes these restrictions and requirements such as:

- Review DB2 10.5 standard parameter settings as indicated in SAP Note 1851832.
- Apply the minimum support packages as indicated in the *SAP Note 1819734*, and listed in Figure 2-2.
- Apply *SAP note 1911087* to fix a known issue and *SAP note 1365982* to update DB6UPDATE and db6_update_db/client scripts.

There is also *SAP note 1825340*, which indicates additional steps and notes for enabling DB2 BLU Acceleration features.

Also, additional recommendations must be verified in the SAP database administration guide *SAP NetWeaver Business Warehouse - Administration Tasks: IBM DB2 for Linux, UNIX, and Windows*[2], at the following website:

http://service.sap.com/~sapidb/011000358700000815432010E

An overview about SAP BW with IBM DB2 for Linux, UNIX, and Windows, with other useful recommendations can be read at SAP Community Network, at the following website:

http://scn.sap.com/docs/DOC-8801

Table 2-2 SAP notes with prerequisites and actions for DB2 BLU Acceleration activation

SAP note	Description
1365982	DB6: Current "db6_update_db/db6_update_client" script (V46). Describes how to update the tools for data conversion to column-organized tables.
1819734	DB6: Use of BLU Acceleration. Describes prerequisites and restrictions for using BLU Acceleration in the SAP environment (includes corrections).
1825340	DB6: Use of BLU Acceleration with SAP BW and applications based on SAP BW. Describes system preparation for DB2 BLU Acceleration activation.

[1] SAP Notes can be found at: http://service.sap.com/notes
[2] SAP DB Administration guide can be found at: http://service.sap.com/instguidesnw

SAP note	Description
1851832	DB6: DB2 10.5 standard parameter settings. Describes prerequisites and the procedure for enabling BLU Acceleration in SAP BW.
1911087	DB6: DROP INDEX to DROP CONSTRAINT. Addresses issues on index creation and constraints.

2.4.1 Software requirements

BLU Acceleration is supported since DB2 10.5 FP3aSAP and SAP BW 7.00, Support Package 31 for some BW objects. Depending on the SAP version and the support package level, other objects may be supported. Figure 2-3, based on *SAP Note 1825340*, summarizes the objects supported depending on the DB2 and SAP BW versions and support package level.

SAP BW Release	DB2 10.5 FP3aSAP			DB2 10.5 FP4 (Cancun)	
	Min	Desired	Supported BW Objects	Desired	Supported BW Objects
7.00	SP31	SP32	- Cumulative and non-cumulative standard InfoCubes: Fact and dimension tables Aggregates The InfoCubes can be part of MultiProviders. Semantically partitioned InfoCubes are also supported. - BW temporary tables Reference: SAP Note #1825340 - DB6: Use of BLU Acceleration with and applications based on	SP33	- Persistent Staging Area (PSA) - Standard DataStore objects - Write-optimized DataStore objects - InfoObjects: Characteristics, time characteristics, units - Flat InfoCube (only for 7.40 SP08 and higher) Reference: SAP Note #1997314 - DB6: Enablement of BLU Acceleration for PSA, DSOs, and Characteristics InfoObjects
7.01	SP14	SP15		SP16	
7.02	SP14	SP15		SP16	
7.11	SP12	SP13		SP14	
7.30	SP10	SP11		SP12	
7.31	SP09	SP11		SP13	
7.40	SP05	SP06		SP07	

Figure 2-3 SAP BW requirements for DB2 BLU

2.4.2 SAP DB2 10.5 bundle and SAP/IBM DB2 license

This section describes the SAP DB2 10.5 bundle and SAP/IBM DB2 license requirements.

SAP 816773 SAP OEM license

Outside the SAP world, DB2 10.5 for Linux, UNIX, and Windows has multiple product editions. The BLU Acceleration feature is supported by the following editions:

- Advanced Enterprise Server (AESE).
- Advanced Workgroup Server (AWS).

You can use BLU Acceleration with the following edition for non-production systems:

- Developer Edition (DE).

Note: If you are referring to an SAP BW System, the *SAP note 816773* must be considered.

Consider, if the DB2 license was bought by SAP (known as an OEM license or ASL license), you can obtain and install the DB2 license files from the SAP Service Marketplace. Otherwise, it is necessary to use the licenses that are received from IBM in the SAP System because you are not allowed to use the license files that are available for download in the SAP Service Marketplace.

SAP note 1260217 explains the components present in the SAP licenses for DB2 for Linux, UNIX, and Windows depending on the version. It may be useful when choosing an upgrade to

decide on the features to be enabled, or if the wanted feature, for example the BLU Acceleration, is available. Refer to Table 2-3.

Table 2-3 SAP notes with information about licensing, components, and features

SAP note	Description
815773	DB6: Installing an SAP OEM license. Describes how to obtain an SAP license and install it.
1260217	DB6: Software components contained in DB2 license from SAP. Describes features present in several DB2 versions for SAP.
1555903	DB6: Supported DB2 database features. Describes DB2 features that can be used in SAP, making reference to relevant notes for such features.

2.4.3 Restrictions for using BLU Acceleration

As mentioned in *SAP Note 1555903 - DB6: Supported DB2 Database Features*, the use of DB2 column-organized tables is restricted to specific usage scenarios. *SAP Note 1819734 - DB6: Use of BLU Acceleration* explains these restrictions further.

- Restrictions to column-organized tables:
  - Conversion of row-organized tables to column-organized tables can be done online. A conversion from column-organized tables to row-organized tables, however, can only be done in read-only mode.
  - With DB2 BLU Acceleration, parallelism at the database level is significantly increased. To avoid over-commitment of CPU resources, you must set a threshold for the number of complex queries that can run at the same time.
- The use of DB2 column-organized tables is restricted to the following scenarios:
  - SAP BW and applications that are based on SAP BW.
  - DB2 Near-Line Storage for SAP BW.
  - The use with other SAP applications is not supported.

2.4.4 General sizing approach for DB2 BLU

The minimum recommendation for a production system, which is the minimum performance consideration, can be extended based on the raw data size that is used in an SAP Business Warehouse production system.

The SAP Quick Sizer calculates the size of the database from the estimated raw data, which consists of the number of rows in InfoCube fact tables, and the active tables of Data Store objects. A general rule, based on several experiences, for database size for SAP BW with DB2 is about a factor 4.7 larger than the raw data including the space for the database transaction logs and the temporary table spaces, plus to PSA, InfoCubes, Data Store, Indexes, and MasterData.

For the minimum performance requirements, refer to the sizing distribution based on raw data for small, medium, and large systems in Table 2-4 on page 23.

Table 2-4 Minimum requirements for performance

	Small	Medium	Large
Raw data size	1 TB	5 TB	10 TB
Cores	8	16	32
Memory	64 GB	256 GB	512 GB

For high-end performance, follow the recommendations as shown in Table 2-5 for the production system.

Table 2-5 High-end requirements for performance

	Small	Medium	Large
Raw data size	1 TB	5 TB	10 TB
CPU	16	32	64
Memory	128 - 256 GB	384 - 512 GB	1024 - 2048 GB

Sizing for DB2 BLU when upgrading from IBM DB2

Using a rule of thumb[3], the sizing method for IBM DB2 BLU database, which source is a DB2 database, considers as input the CPU and memory and applies the following formulas:

```
CPU (BLU) = CPU * 1,3
Memory (BLU) = Memory * 2
```

This sizing must consider too, the minimum recommendation of 8 cores and 64 GB RAM.

2.4.5 Storage requirements

There is no special storage requirement for DB2 with BLU Acceleration. A general suggestion is to store column-organized tables on storage systems with good random read and write I/O performance. This notably helps in situations where active table data exceeds the main memory available to DB2 or when temporary tables are populated during query processing.

In these situations, the following storage types can be of benefit:

- Solid-state drives (SSDs).
- Enterprise SAN storage systems.
- Flash-based storage with write cache.

The following IBM storage systems exhibit good random I/O characteristics:

- IBM FlashSystem™ 810, 820, and 840.
- IBM Storwize® V7000 (loaded with SSDs).
- IBM XIV® storage system.
- IBM DS8000® Series.

For more information about DB2 BLU Architecture, see *Architecting and Deploying DB2 with BLU Acceleration*, SG24-8212-01:

http://www.redbooks.ibm.com/abstracts/sg248212.html

[3] Rule of thumb (RoT) is always based on best guess and assumptions estimated on a running balanced system, with the purpose of giving a basis to discuss with the customer.

2.4.6 Platform requirements

BLU Acceleration is supported on AIX and Linux x86 (Intel and AMD) platforms using 64-bit hardware as shown in Table 2-6. To take advantage of BLU Acceleration fast analytics, the following platforms are suggested:

- IBM POWER7® or later.
- Intel Nehalem (or equivalent) or later.

Table 2-6 Platform requirements

Operating system	Operating system minimum version	Hardware recommendation
AIX	AIX 6.1 TL7 SP6 AIX 7.1 TL1 SP6	IBM POWER7 or later
Linux x86 64-bit	RHEL 6 SLES 10 SP4 SLES 11 SP2	Intel Nehalem (or equivalent) or later

For the most up-to-date information about DB2 system requirements, refer to the following website:

http://www-01.ibm.com/support/docview.wss?uid=swg27038033#105

For more information about IBM Power Architecture®, see Chapter 3, “Optimizing IBM Power Systems for SAP BW using IBM DB2 BLU” on page 25.

3

Optimizing IBM Power Systems for SAP BW using IBM DB2 BLU

No tuning is required to use DB2 BLU Acceleration. In general, users only need to follow the BLU best practices for capacity planning, and recommended configuration as documented in the previous chapters. The default configuration, in most cases, produces good analytical results.

However, depending on the customer requirements and workload characteristics, some considerations and tuning are required to achieve optimal results to help meet the business objectives.

This chapter provides information to optimize IBM Power Systems for SAP BW using IBM DB2 BLU Acceleration.

The following topics are covered:

- TurboCore mode of the IBM POWER7
- Start the most important LPAR for memory affinity
- Set DB2_RESOURCE_POLICY for memory affinity
- Dedicated logical partition for stable performance
- Active memory expansion is not recommended in this case
- AIX tunables
- Sizing approach with IBM Power Systems

© Copyright IBM Corp. 2015. All rights reserved.

3.1 TurboCore mode of the IBM POWER7

TurboCore mode in IBM POWER7 increases processor frequency, and increases the L3 cache per core. Enabling TurboCore mode in POWER7 is useful because this processing mode helps achieve higher BLU performance.

TurboCore is a special processing mode as four of the eight cores and L2 cache are turned off as shown in Figure 3-1.

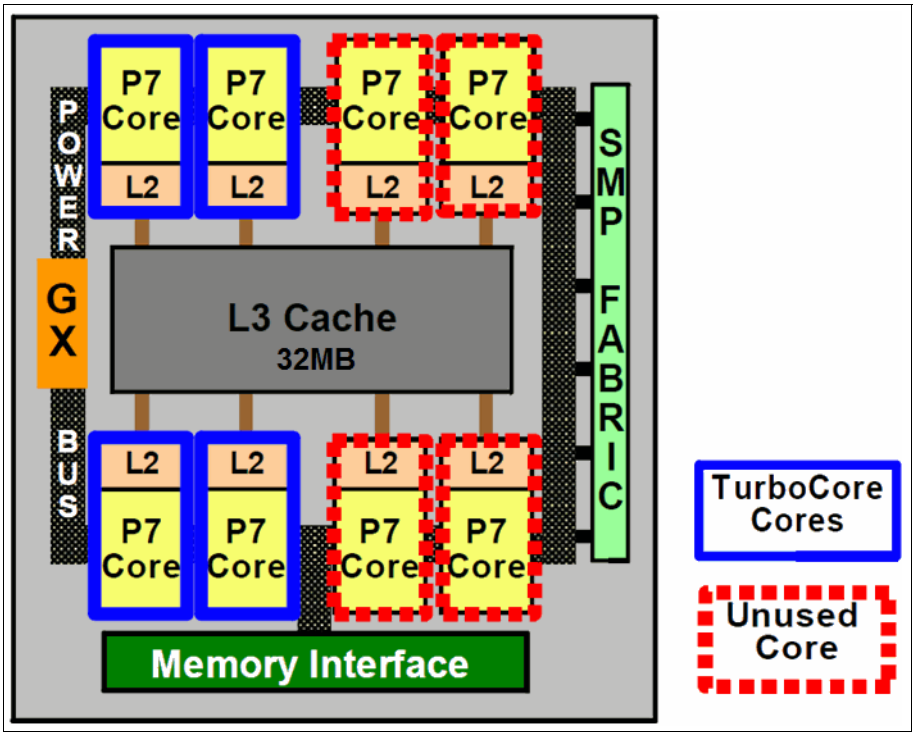

Figure 3-1 POWER7 TurboCore mode

By enabling TurboCore mode, the processor cores execute at a higher frequency and the L3 cache per core increases 4 - 8 MB. In TurboCore mode, the CPU frequency increases as shown in Table 3-1.

Both, the higher frequency and the larger amount of cache per core provide better performance for an application that requires more cache, for example the IBM DB2 BLU Acceleration.

Table 3-1 CPU frequency in TurboCore mode

TurboCore mode supported model	Normal mode	TurboCore mode
IBM Power 795 server	4.00 GHz	4.25 GHz
IBM Power 780 server	3.92 GHz	4.14 GHz

Note: TurboCore mode is no longer supported on the IBM Power 780 server.

You can enable TurboCore mode with the Advanced System Management Interface (ASMI) as shown in Figure 3-2.

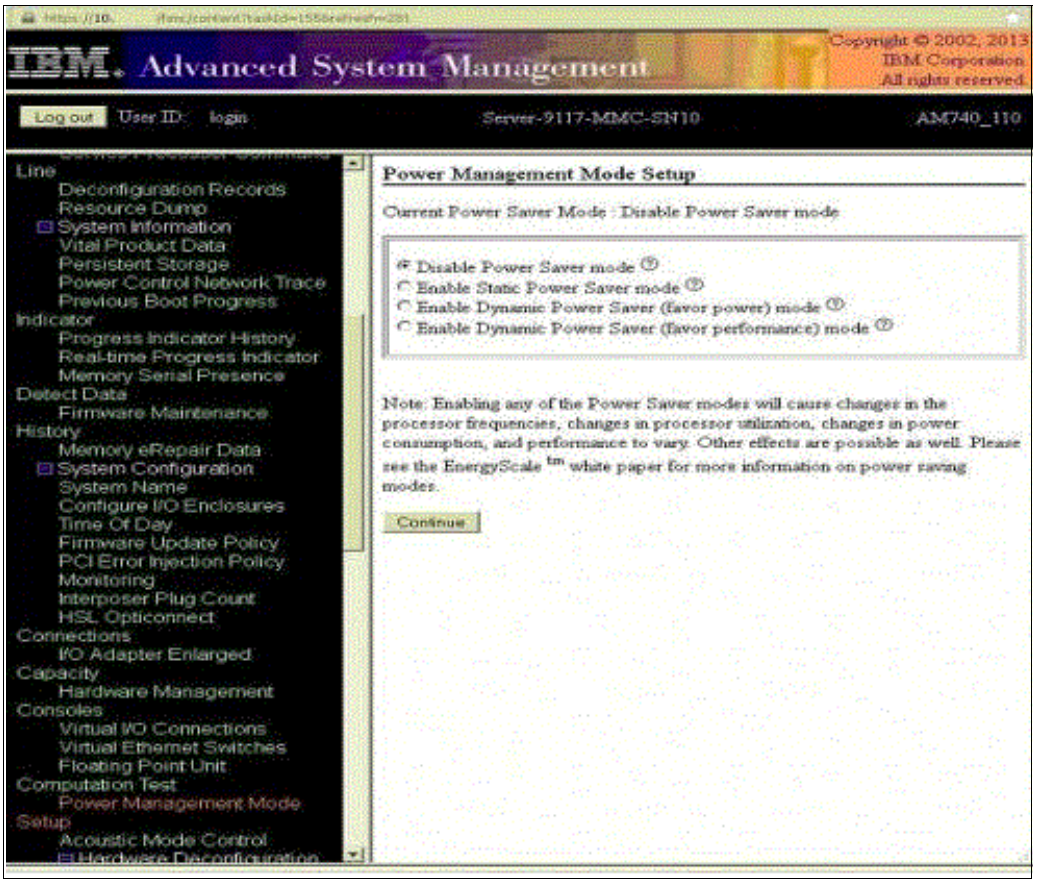

Figure 3-2 ASMI menu to enable TurboCore mode

Note: You can only switch TurboCore mode on or off for the entire system.

3.2 Start the most important LPAR for memory affinity

The order of starting the dedicated logical partitions (LPARs) is one consideration to keep in mind while looking for obtaining the best performance. When you have multiple partitions in a POWER system, start the most important partition first.

Before providing the details on the order of starting the partitions, take a quick look at the Non-Uniform Memory Access (NUMA) architecture. NUMA is a multi-processor architecture, and the IBM POWER processor is based on NUMA architecture. In NUMA, memory access time depends on the memory location relative to a processor. The architecture design of the POWER platform is mostly NUMA with three levels (Figure 3-3 on page 28):

- Each POWER7 chip has its own memory dimms. Access to these dimms has a low latency and is named *local memory access*.

- Up to four POWER7 chips can be connected to each other in the same CEC (or node) by using XYZ memory buses. Access to memory owned by another POWER7 chip in the same CEC is called *near* or *remote memory access*. Near or remote memory access has a higher latency compared to local memory access.
- Up to eight CECs can be connected through AB memory buses from a POWER7 chip (only on high-end systems). Access to memory owned by another POWER7 in another CEC (or node) is called *far* or *distant memory access*. Far or distant memory access has a higher latency than remote memory access.

These memory access latency concepts also apply to the IBM POWER8™.

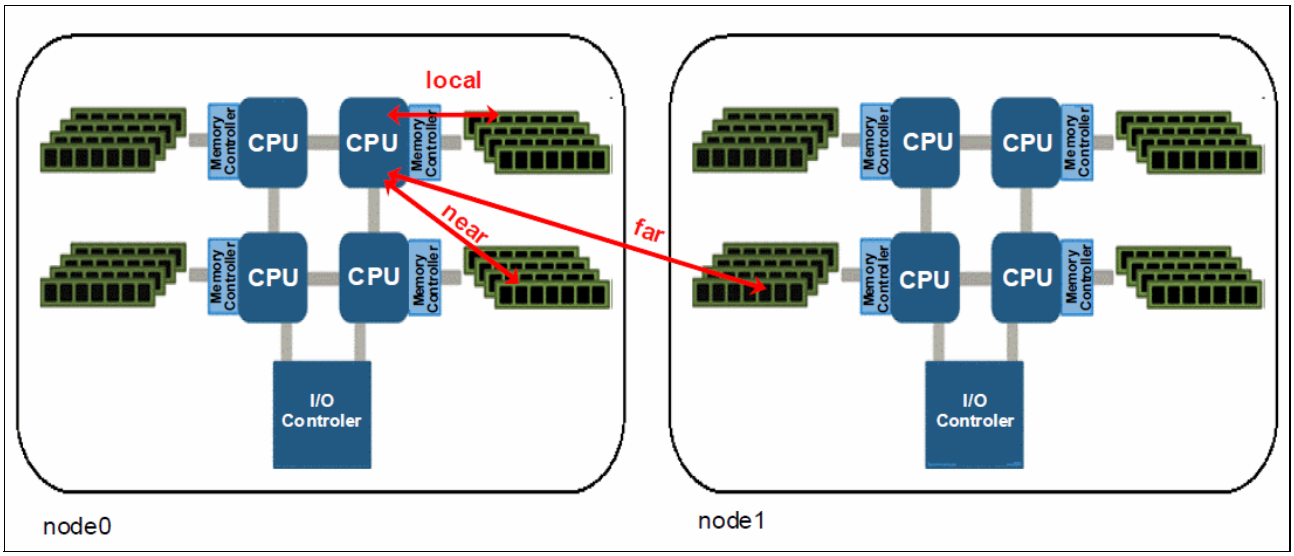

Figure 3-3 POWER7 memory access latency concepts

With larger dedicated partitions, for example partitions using DB2 BLU Acceleration, one of the most important performance considerations is which cores and memory are allocated to a partition. The IBM POWER Hypervisor™ attempts to allocate all cores for the partition from a single POWER chip and attempts to allocate only memory local to the chip. These partitions generally and automatically have good cache and memory affinity.

However, it might not be possible to obtain resources for each of the LPARs from a single chip. For example, assume that you have a 32-core system with four chips, each with eight cores. If five partitions are configured, each with six cores, the fifth LPAR spreads across three chips. Hence, start the most important partition first to obtain resources from a single chip.

This consideration for the order of starting the partitions is only necessary when you install a new system or you turn on the POWER system. After a partition has been powered on, the server has been restarted so that the order of activation is not important. On the Hardware Management Console (HMC), there is an option to activate the current configuration as shown in Figure 3-4 on page 29, and when you use this option, there is no change in the current placement of the partition. Although, activating with a partition profile might change the current placement of the partition.

For more information about the POWER processor, see *Performance Optimization and Tuning Techniques for IBM Processors, including IBM POWER8*, SG24-8171-00:

http://www.redbooks.ibm.com/abstracts/sg248171.html

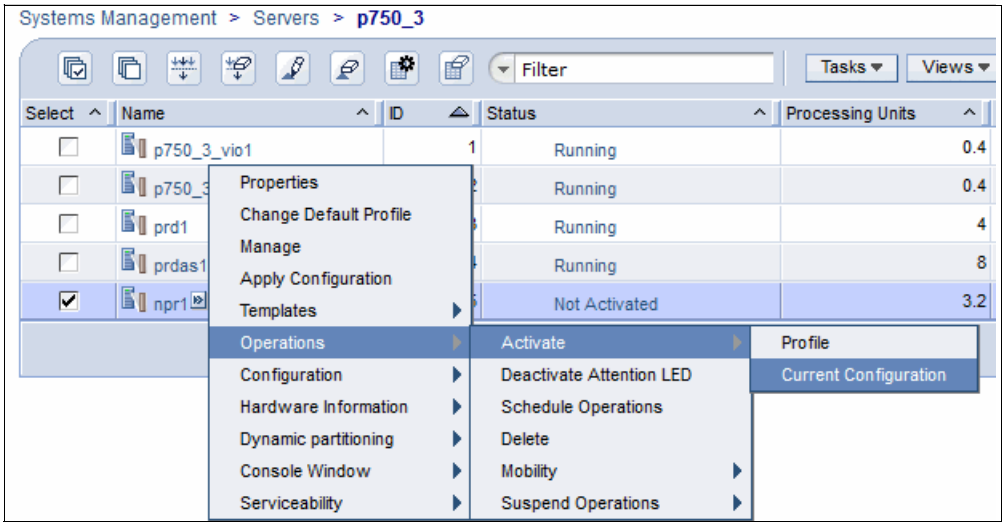

Figure 3-4 Activate the LPAR using the current configuration option

3.3 Set DB2_RESOURCE_POLICY for memory affinity

Setting the DB2 registry variable **DB2_RESOURCE_POLICY** to AUTOMATIC is also an effective way to minimize the NUMA effect if your DB2 version is 10.5 FP4 or above.

This variable defines a resource policy that can be used to limit what operating system resources are used by the DB2 database. The variable contains rules for assigning specific operating system resources to specific DB2 database objects. For example, this registry variable can be used to limit the set of processors that the DB2 database system uses.

In the NUMA architecture, binding each individual DB2 process to a particular resource set can be beneficial in performance tuning scenarios.

For IBM POWER7 Systems™ running AIX 6.1 Technology Level (TL) 5 or higher, this variable can be set to AUTOMATIC. This feature is a new one introduced in DB2 10.5 FP4 and can improve workload performance.

With this registry setting enabled, the DB2 database system automatically determines the hardware topology and assigns engine dispatchable units (EDUs) to the various hardware modules in such a way that memory can be more efficiently shared between multiple EDUs that must access the same regions of memory. The AUTOMATIC setting also determines whether to enable memory affinitization, whereby EDUs attempt to allocate local memory during processing.

This registry setting is intended for larger POWER7 Systems with sixteen or more cores, which can help enhance query performance for some workloads.

Example 3-1 shows how to set the registry variable to AUTOMATIC.

Example 3-1 Setting the registry variable to AUTOMATIC

```
$ db2set DB2_RESOURCE_POLICY=AUTOMATIC
$ db2start
```

For more information about the **DB2_RESOURCE_POLICY**, see “DB2 Version 10.5” of the IBM Knowledge Center, in the performance variables section, which describes the DB2 performance registry variables, including the **DB2_RESOURCE_POLICY**:

http://ibm.co/1M5xWWl

3.4 Dedicated logical partition for stable performance

Dedicated processor logical partitions (LPARs) provide constant stability and the best performance for DB2 BLU Acceleration without being affected by the workload of another LPAR in the same physical server.

A logical partition (LPAR) is an isolated computing domain. Each LPAR has its own processing capacity, memory, network connection, storage, and operating system. There are two types of partitions in a POWER system: Dedicated LPAR and shared LPAR. Figure 3-5 shows a sample configuration with a dedicated LPAR and shared LPARs.

The left-most partition shows a dedicated LPAR having two physical processors exclusively. If you use a dedicated LPAR, you must assign at least one processor to the dedicated LPAR.

The other partitions show shared LPARs. The hypervisor dispatches the processing capacity to the LPARs from a shared processor pool. Processing capacity can be configured with units of 1/100 of a processor. The minimum amount of processing capacity that has to be assigned to a partition is 1/20 of a processor (latest POWER system firmware levels).

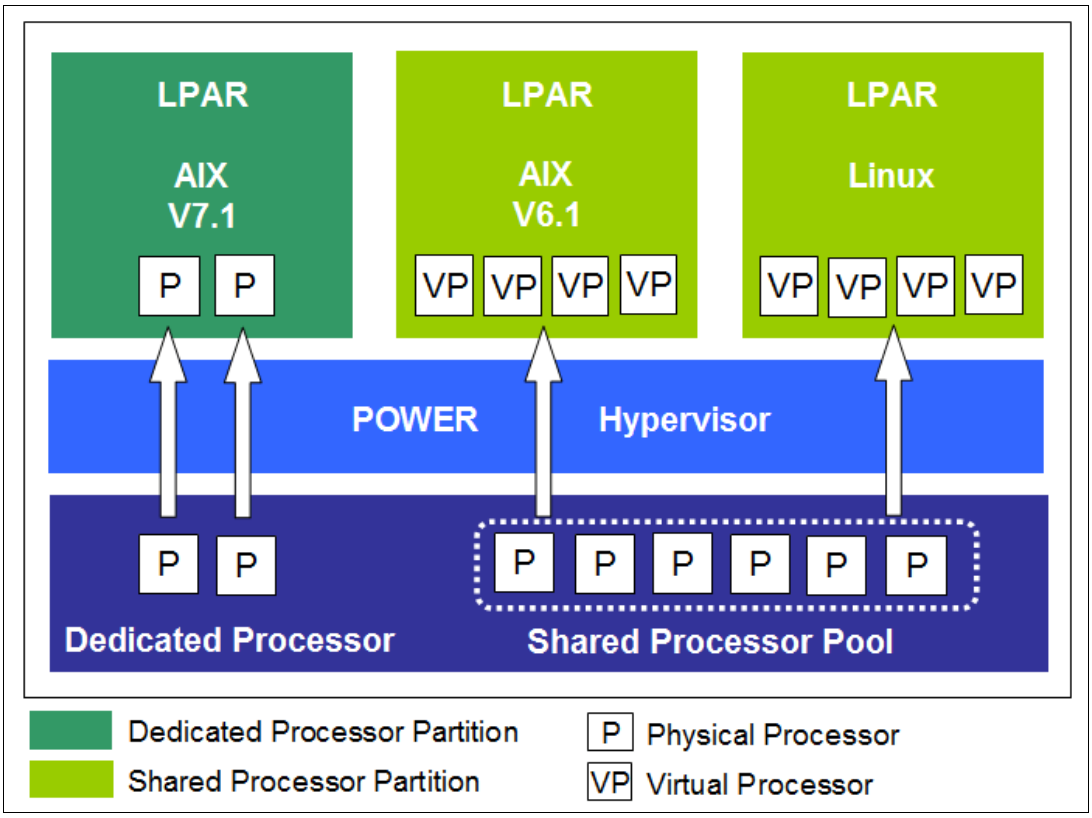

Figure 3-5 Dedicated LPAR and shared LPARs

DB2 BLU Acceleration is supported on dedicated LPARs and shared LPARs. It is important to consider the characteristics of these LPAR types.

The biggest benefit of dedicated processor LPARs is that there is no impact due to another LPAR's workload because the dedicated LPARs hold processors for exclusive use. If you use IBM POWER6® based and later servers, you can design a flexible system using a feature called *shared dedicated capacity*. This feature offers the capability of harvesting unused processor cycles from dedicated-processor partitions. These unused cycles are then donated to the physical shared-processor pool associated with IBM Micro-Partitioning®. This ensures the opportunity for maximum processor utilization throughout the system.

The biggest benefit of shared processor LPARs is a more flexible configuration. Shared processor LPARs have a specific processing mode that determines the maximum processing capacity that is given to them from their shared-processor pool. One mode is the uncapped mode and it enables the processing capacity to exceed automatically the entitled capacity when resources are available in their shared-processor pool. Another mode is the capped mode and it means the processing capacity that is given can never exceed the entitled capacity of the shared LPAR. These flexible configurations provide high processor utilization for the entire POWER server. However, the shared LPAR is affected by the workload of another LPAR in the same shared-processor pool.

There is no general guide on how much DB2 BLU Acceleration is affected by other LPARs in the same server because it depends on the characteristics of the other applications that are sharing the system.

As reference, we introduce how much analytics queries are affected by other LPAR workloads by using dedicated LPAR and shared LPARs. There are comparative test results of the query elapsed time in three test cases as shown in Figure 3-6 on page 32.

In TESTCASE1, there are two LPARs and two VIOS in one Power 750 machine and DB2 BLU is running on a dedicated LPAR (LPAR2) that has dedicated CPU, dedicated memory, and physical Fibre Channel cards. In this test, LPAR1 can be ignored.

In TESTCASE2, BLU is running in a shared LPAR (LPAR2). Another LPAR (LPAR1) is uncapped shared LPAR using the same CPU pool with LPAR1. The CPU pool contains 16 CPUs. In this test case, no operation is running in LPAR1 and there is little CPU utilization on LPAR1. This means LPAR1 does not have any significant impact on DB2 BLU in LPAR2.

In TESTCASE3, the LPAR configuration is the same as with TESTCASE2. Heavy operations are using all available processor capacity exceeding the estimated capacity (EC) on LPAR1. This means LPAR1 has the biggest impact on LPAR2 with DB2 BLU. We executed the stress tool called nmem64[1] on LPAR1 in 16 parallels.

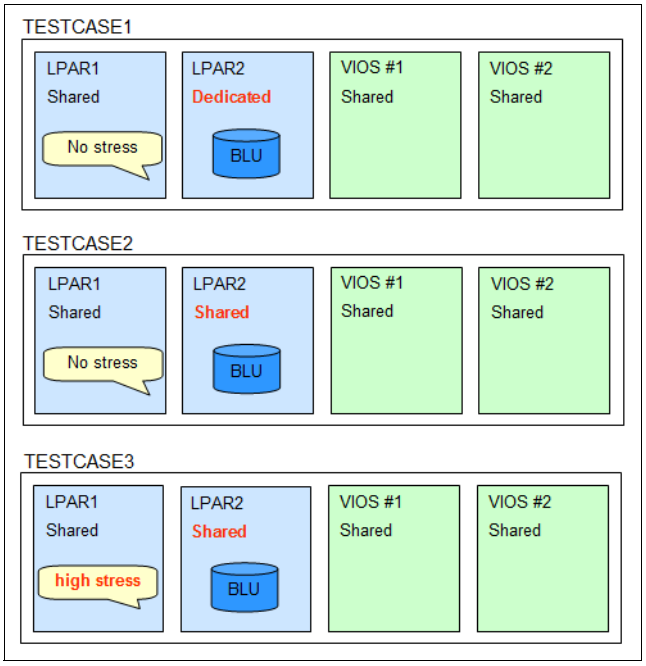

Figure 3-6 Test environment

Detailed information about the LPAR configuration is shown in Figure 3-7.

		LPAR1	LPAR2 (BLU)	VIOS1	VIOS2
TESTCASE 1	dedicate/shared	shared	**dedicate**	shared	shared
	EC	4.0	8	0.4	0.4
	VP	16	–	2	2
	cap mode	uncap	–	uncap	uncap
	MEMORY	28	64	4	4
	Disk IO	NPIV	physical FC card	physical FC card	physical FC card
TESTCASE 2	dedicate/shared	shared	**shared**	shared	shared
	EC	4.0	8.0	0.4	0.4
	VP	16	8	2	2
	cap mode	uncap	cap	uncap	uncap
	MEMORY	28	64	4	4
	Disk IO	NPIV	physical FC card	physical FC card	physical FC card
TESTCASE 3	dedicate/shared	shared	**shared**	shared	shared
	EC	4.0	8.0	0.4	0.4
	VP	16	8	2	2
	cap mode	uncap	cap	uncap	uncap
	MEMORY	28	64	4	4
	Disk IO	NPIV	physical FC card	physical FC card	physical FC card

Figure 3-7 LPAR configuration in the test

[1] See http://ibm.co/1Ce52vt

The elapsed time can be found by executing each query three times, and calculating the average of the second and third elapsed time. The query used in this test is the Star Schema Benchmark (SSB) queries (Q1.1~Q4.3) and some additional queries (test1~test4). In this test, about 60 GB of raw data, which is relatively small for DB2 BLU, is loaded into the DB2 BLU database.

The SSB used in this test is designed to measure the performance of database products in support of classical data warehousing applications. For more details about the SSB benchmarks, see the SSB website:

http://www.cs.umb.edu/~poneil/StarSchemaB.PDF

TESTCASE1 shows in Figure 3-8 and Figure 3-9 that a dedicated LPAR can provide constant stability and the best performance for DB2 BLU Acceleration. TESTCASE3 shows that the analytic query on LPAR2 is affected mostly by LPAR1's workload.

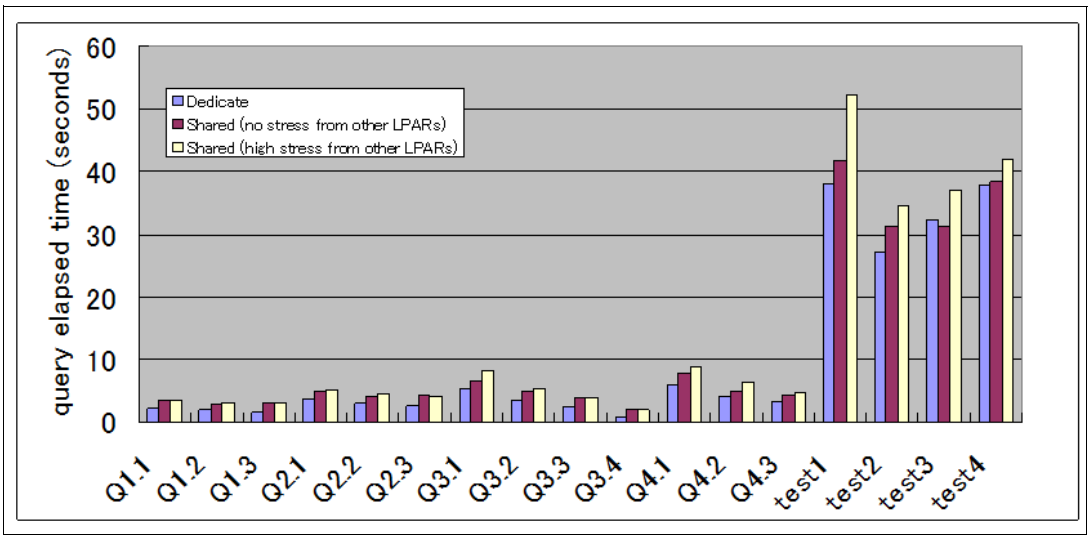

Figure 3-8 Test result

elapsed time (seconds)

	Q1.1	Q1.2	Q1.3	Q2.1	Q2.2	Q2.3	Q3.1	Q3.2	Q3.3	Q3.4	Q4.1	Q4.2	Q4.3	test1	test2	test3	test4
testcase1	2.15	1.91	1.55	3.59	3.13	2.59	5.51	3.46	2.43	0.78	5.78	4.18	3.17	38.1	27	32.3	37.82
testcase2	3.37	2.9	2.99	4.92	4.22	4.23	6.55	5.06	3.86	1.96	7.79	5.05	4.26	41.79	31.18	31.29	38.29
testcase3	3.55	3	3.02	5.25	4.65	4.19	8.3	5.34	3.91	1.91	8.92	6.22	4.85	52.23	34.37	37.08	41.95

Figure 3-9 Test results elapsed time

The preceding performance data that is shown was obtained in a controlled environment. Results that are obtained in your environment might vary significantly. There is no guarantee that the same results will be obtained elsewhere.

If you use DB2 BLU Acceleration in a shared LPAR, the following guides can help:

- Server virtualization with IBM PowerVM®

  http://www-03.ibm.com/systems/power/software/virtualization/resources.html
- IBM POWER7 Virtualization Best Practice Guide

  http://ibm.co/1K5NLbj
- IBM PowerVM Virtualization Introduction and Configuration, SG24-7940-05

  http://www.redbooks.ibm.com/redbooks/pdfs/sg247940.pdf

3.5 Active memory expansion is not recommended in this case

Active memory expansion (AME) is a capability that is supported on POWER7 or later servers that employs memory compression technology to expand the effective memory capacity of an LPAR. By enabling active memory expansion for the LPARs, the operating system recognizes more memory than the available physical memory.

The use of AME for SAP application servers with a DB2 database server is generally supported. However, the use of AME for DB2 database servers is not recommended because of the additional overhead the CPU creates to enable AME becomes an issue when used with DB2.

Before using AME, consider using DB2 static/adaptive compression as well as buffer pools because this reduces the demand on the I/O subsystem and saves disk space. In DB2 BLU Acceleration, data is compressed by default.

For more information about using AME for DB2, see SAP Note 1906519 and the technote listed in the website:

http://www-01.ibm.com/support/docview.wss?uid=swg21567898

3.6 AIX tunables

There is not DB2 BLU-specific AIX tuning. You only need to follow the DB2 recommended values for the AIX tunables. For the DB2 recommended value of AIX tunables, see *Best Practices for DB2 on AIX 6.1 for Power Systems*, SG24-7821-00:

http://www.redbooks.ibm.com/abstracts/sg247821.html

If you have multipath that is enabled for the storage, refer to the following materials as these show how to tune the disk and the Fibre Channel adapters to improve performance:

- IBM AIX MPIO: Best practices and considerations:

  http://www.ibm.com/developerworks/aix/library/au-aix-mpio

- AIX and VIOS Disk and Fibre Channel Adapter Queue Tuning:

  http://www-03.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/TD105745

If you run SAP on AIX, see the SAP Notes 1972803 and the SAP on Power Systems Best Practices Guide. These documents provide recommendations and clarifications for many topics relevant for running SAP on AIX:

- IBM SAP Technical Brief, SAP on Power Systems Best Practices Guide:

  http://ibm.co/1x1QcTi

3.7 Sizing approach with IBM Power Systems

DB2 BLU Acceleration is supported on AIX and Linux x86 (Intel and AMD) platform using 64-bit hardware. To take advantage of DB2 BLU Acceleration for fast analytics, deploy IBM POWER7 or later systems because SIMD was improved in POWER7 and has received further improvements in POWER8.

Figure 3-10 shows the POWER7 model depending on the row data size.

ROW data Size, Uncompressed (TB)	System Type	# of Cores	RAM (GB)
<1TB	710 or higher	8 -12	64 - 128
1 – 2.5	730 or higher	8 - 16	128 - 256
2.5 – 5	730 or higher	16 - 24	256 - 512
5 – 7.5	740 or higher	16 - 32	512 - 1024
7.5 – 10	750 or higher	32 - 48	1024 - 2048
10 – 15	760 or higher	48 - 64	1024 - 2048
15 – 20	770 or higher	64 - 128	2048 - 4096
20 +	770 or higher	64 +	2048 +

Figure 3-10 POWER7 model depends on row data size

When it comes to SAP environments, SAP products are supported on IBM POWER hardware, which is supported by IBM for the AIX versions that are documented in the Product Availability Matrix (PAM) available at the following site:

http://www.service.sap.com/pam

4

SAP BW landscape scenario with IBM DB2 BLU Acceleration

This chapter describes a common scenario: How to build up an SAP NetWeaver Business Warehouse (known as SAP NetWeaver BW, or SAP BW) landscape with IBM DB2 for Linux, UNIX, and Windows with BLU Acceleration and the Near-Line Storage (NLS) feature on Power Systems.

There are some challenges in order to fulfill all BLU requirements mentioned in Chapter 2, “Requirements for an SAP NetWeaver BW environment” on page 17. Therefore, this scenario serves as a template on how to achieve the target scenario in the following customer situations:

1. Customer without SAP BW (new system installation)
2. Customer with SAP BW but on other platform or database (SAP migration)
3. Customer with SAP BW on POWER and DB2, but not BLU enabled (upgrades)

The following topics are discussed in this chapter:

- Design SAP BW landscape
- Implementation of SAP BW with DB2 BLU
- Administration and operation with SAP DB2 BLU

© Copyright IBM Corp. 2015. All rights reserved.

4.1 Design SAP BW landscape

This environment consists of three systems, which reflect a common SAP Business Warehouse System landscape to assure SAP software logistics.[1] The SAP identifier for each system is shown in Figure 4-1.

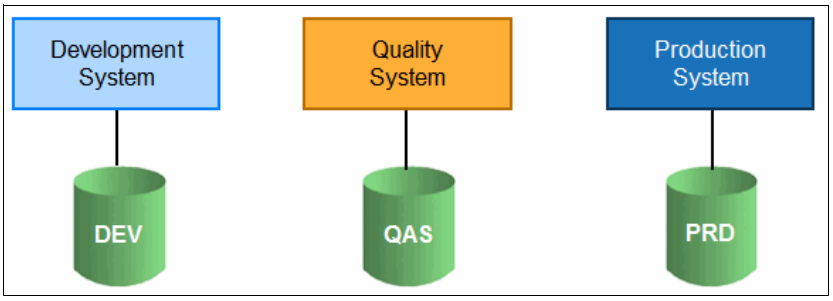

Figure 4-1 SAP system landscape

As shown in Chapter 1, “Introducing SAP NetWeaver Business Warehouse with IBM DB2 BLU” on page 1, the result of the proof of concept (POC) with DB2 BLU demonstrates that the BLU Acceleration of your SAP BW System can supersede the implementation of the SAP Business Warehouse Accelerator (BWA). Therefore, this design scenario is without BWA.

Note: Because of the BLU improvements on SAP BW InfoProvider, there is no need to copy InfoCubes and DSO to a BWA. However, if you want to use SAP BW workspaces, you need to establish a connection to BWA or SAP Hana.[a]

a. See http://help.sap.com/businessobject/product_guides/AMS14/en/14_aaoffice_user_en.pdf, on page 27.

4.1.1 Hardware resources for SAP BW landscape with DB2 BLU

This SAP Business Warehouse landscape scenario is sized with the minimum requirements, which are listed in *SAP Note 1851832 (DB6 DB2 10.5 standard parameter settings)*:

- Development system has at least 4 CPUs and 32 GB.
- Quality assurance system has at least 4 CPUs and 32 GB.
- Production system has at least 8 CPUs and 64 GB.

Note: These are the minimum requirements for DB2 BLU to use BLU Acceleration. For more details about sizing approaches, see 2.4.4, “General sizing approach for DB2 BLU” on page 22 and 3.7, “Sizing approach with IBM Power Systems” on page 34.

The baselines are listed in Figure 4-2.

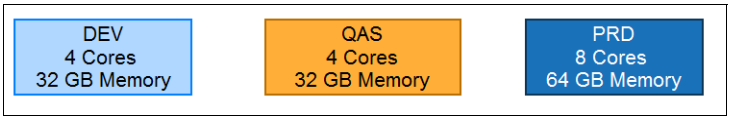

Figure 4-2 CPU and memory baselines for the SAP BW System landscape

[1] See http://bit.ly/1M5FPep

The hardware environment consists of two IBM Power 750 servers, a Hardware Management Console (HMC), and IBM DS8000 storage as shown in Figure 4-3. The IBM Power 750 server is configured with IBM POWER7+™ 3.3 GHz 16cores and 128 GB RAM. The IBM DS8000 is connected to the Power servers via the storage area network (SAN).

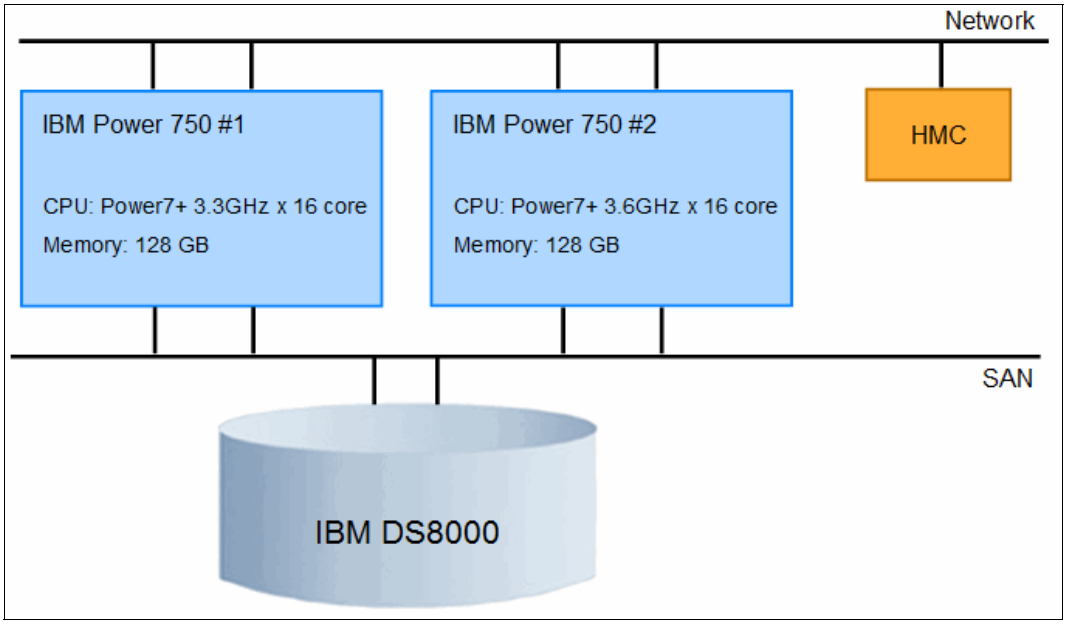

Figure 4-3 Hardware environment for SAP BW using BLU Acceleration

The entire systems landscape and the NLS implementation should not be on one box (physical server) for high availability reasons. If there is a failure, the whole landscape is not available. Therefore, in this example there are two servers: One server is for preproduction systems, and a second server is for the production systems:

- p750#1 is for the preproduction system consisting of DEV and PRD.
- p750#2 is for the production system PRD.

Separating the servers also helps to improve security and reduces effects from preproduction workload on the preproduction system and production system. This general design allows to add high availability features such as IBM PowerHA or HADR to the production systems.

4.1.2 Considerations for the preproduction systems (DEV and QAS)

The development system (DEV) is usually a small configuration because a few developers use it, and there is only some test data available. The Quality System (QAS) is mainly a medium-sized system. It contains more test data, but only some users are testing the new data. Performance is not the main driving factor for preproduction systems, so that the BLU operations can take the necessary CPU time out of the CPU pool and use the VIOs. The CPU resources can be in uncapped mode.

The installation for DEV and QAS are standard system installations. This means that the SAP Primary Application Server (PAS), the SAP Central Service Instances (ASCS), and the DB2 for Linux, UNIX, and Windows database are installed on one host (LPAR).

The preproduction server (box) focuses on the flexible utilization of the resources. The Power750#1 box consists of three shared LPARs and two VIOS as shown Figure 4-4.

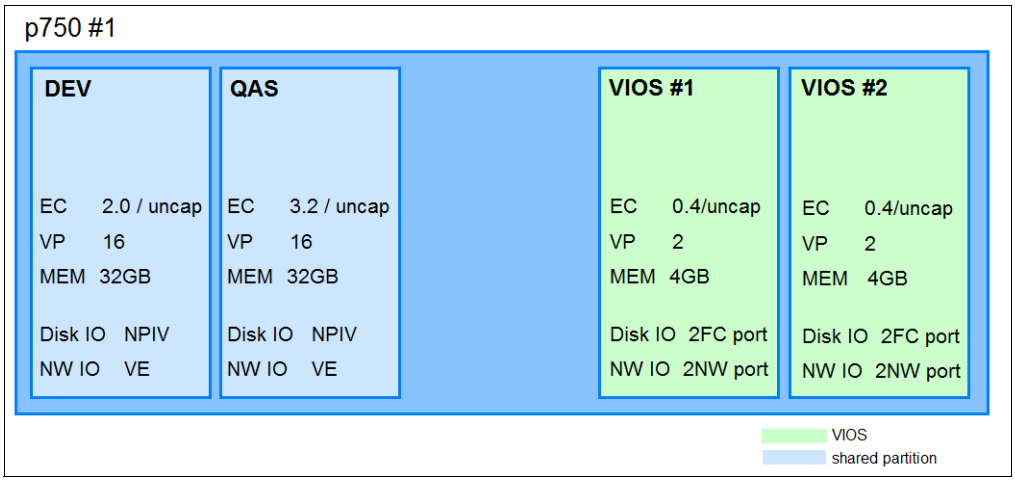

Figure 4-4 LPAR configuration for the preproduction box

For the in-network I/O and for the storage I/O configurations on each LPAR, multiple routes are provided via the dual VIOS using Virtual Ethernet (VE) and N_Port ID Virtualization (NPIV).

For more information about the VIOS, NPIV, and VE features, refer to the following IBM Redbooks publications:

- IBM PowerVM Best Practices, SG24-8062:

  http://www.redbooks.ibm.com/redbooks/pdfs/sg248062.pdf

- IBM System p® Advanced POWER Virtualization Best Practices, REDP4194:

  http://www.redbooks.ibm.com/redpapers/pdfs/redp4194.pdf

4.1.3 Consideration for the production system (PRD)

The production system, first and foremost, has to deliver performance to the user and has to manage big data. To guarantee BLU performance and good usage of CPU caches, the production system should be set up in a dedicated POWER environment, so that CPU and I/O resources are available exclusive for BLU.

You can decide the location of the SAP Central service instance for ABAP (ASCS instance), the database instance (DB), and the primary application server instance (PAS) with the installation options. SAP provides the following installation options: Standard, distributed, and HA installation options. With the standard system installation, you can install all SAP instances on a single host. With a distributed system installation, you can distribute each SAP instance to separate hosts.

In this scenario, we designed a distributed SAP BW installation for PRD. This means that SAP PAS and SCS operate on a host that is separated from the database server. Both SAP instances are still running in a shared environment for flexible utilization of resources, but the DB2 database runs in a dedicated environment. This approach provides the possibility to run the SAP environment on a flexible and shared POWER environment.

As shown in 3.4, “Dedicated logical partition for stable performance” on page 30, the performance results for a dedicated BLU environment are better than on shared LPAR. The

performance for a shared LPAR is also good, but you face some overhead due to the workload of the hypervisor managing the shared environment. Also, the CPU cache is not always guaranteed to BLU.

The production server focuses on both performance and flexible utilization of resources. The Power750#2 server consists of two shared LPARs, one dedicated LPAR, and two VIOS, as shown Figure 4-5.

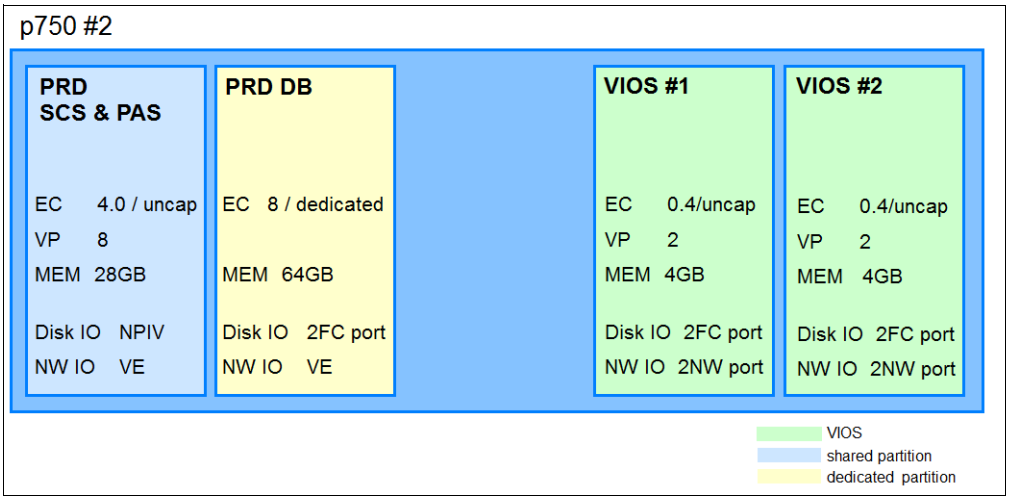

Figure 4-5 LPAR configuration for the production server

HA solutions are not in scope in this scenario. If you want more information about HA solutions in an SAP environment, refer to the following SAP IBM AIX High Availability, Disaster Recovery and Storage website:

http://scn.sap.com/docs/DOC-8761

4.1.4 DB2 database layout for BLU

A default SAP BW installation creates a default single partition without intrapartition. The layout of your database has a large impact on the performance of your SAP BW System and requires careful planning.[2]

For a regular SAP system with DB2 but without BLU, you have the following database layout options that are available:

- Single partition with intrapartition parallelism.
- Physical or logical database partitioning feature (DPF).
- IBM DB2 pureScale® feature.

BLU is only supported in a single partition with intrapartition parallelism.

Note: Column-organized tables (BLU) cannot be combined with the database partitioning feature or pureScale. If you use DPF or DB2 pureScale and want to use BLU, first you have to unpartition your DB2 database or leave the DB2 pureScale environment. But the DB2 HADR solution including BLU Acceleration is supported since the DB2 Cancun release v10.5 FP4.

[2] See http://service.sap.com/~sapidb/011000358700001420572010E (login required)

4.1.5 Migrate SAP BW to Power Systems with AIX

In the following situations, for an existing SAP BW environment, an SAP Unicode conversion, an SAP migration, or at least an SAP upgrade is mandatory:

- SAP BW System is Non-Unicode.
- SAP BW System is not on a supported platform (for example HP-UX).
- SAP BW System without DB2 database (for example, Oracle database).
- SAP BW System with an old SAP BW release that is not BLU enabled.
- SAP BW System with DB2 DPF feature enabled.

For information about SAP DB2 Migrations and Unicode Conversions, refer to *DB2 Optimization Techniques for SAP Database Migration And Unicode Conversion*, SG24-7774:

http://www.redbooks.ibm.com/abstracts/sg247774.html?Open

In some circumstances, the need to do a combined upgrade and Unicode conversion should be considered. For information about designing an IBM Power Systems environment for SAP BW, refer to 4.1, “Design SAP BW landscape” on page 38.

4.2 Implementation of SAP BW with DB2 BLU

The SAP ABAP installations are based on Release 7.40 SR2 with DB2 for Linux, UNIX, and Windows Version 10.5 FP4. The entire procedure on how to perform an SAP NetWeaver system installation is described in the SAP installation guide, which is available in the SAP Service Marketplace.[3]

This section describes three important points that are either mandatory or a highly recommended best practice for DB2 BLU Acceleration:

- Adaptive row compression: Highly recommended.
- Automatic storage management (ASM): Mandatory.
- Multi-temperature storage groups: Recommended.

4.2.1 Enabling adaptive row compression

In an SAP BW environment, column and row organized tables do coexist. As described in “Column store” on page 6, column-organized tables are always in a compressed state. However, row-organized tables per default are not. To maximize disk space savings, it is recommended to enable data compression also for row-organized tables during an SAP installation or SAP import. This can be done easily as shown in Figure 4-6 on page 43. If you do not choose the data compression option, tables are created *without* adaptive data compression. Although, new tables are automatically created with DB2 data compression because the global compression option is activated per default with SAP BW Release 7.40, as described in *SAP Note 1690077* (*DB6: Global compression option),* and as shown in Example 4-1.

Example 4-1 Global compression option for newly created tables in DEV

```
db2 "VALUES( SAPDEV.GLOBAL_COMPRESSION_OPTION )"

1
-----------------------------------------------
YES
```

[3] See http://service.sap.com/~sapidb/012002523100011883442014E (logon required)

Furthermore, there is no need to set the DB2 profile DB2_ROWCOMPMODE_DEFAULT= ADAPTIVE anymore because with DB2 v10.5, it is included in the DB2 registry DB2_WORKLOAD=SAP per default. Additionally online inplace REORG is supported on row tables with adaptive compression with v10.5.

Note: The following screen capture is an example of the integration between IBM DB2 BLU and SAP BW.

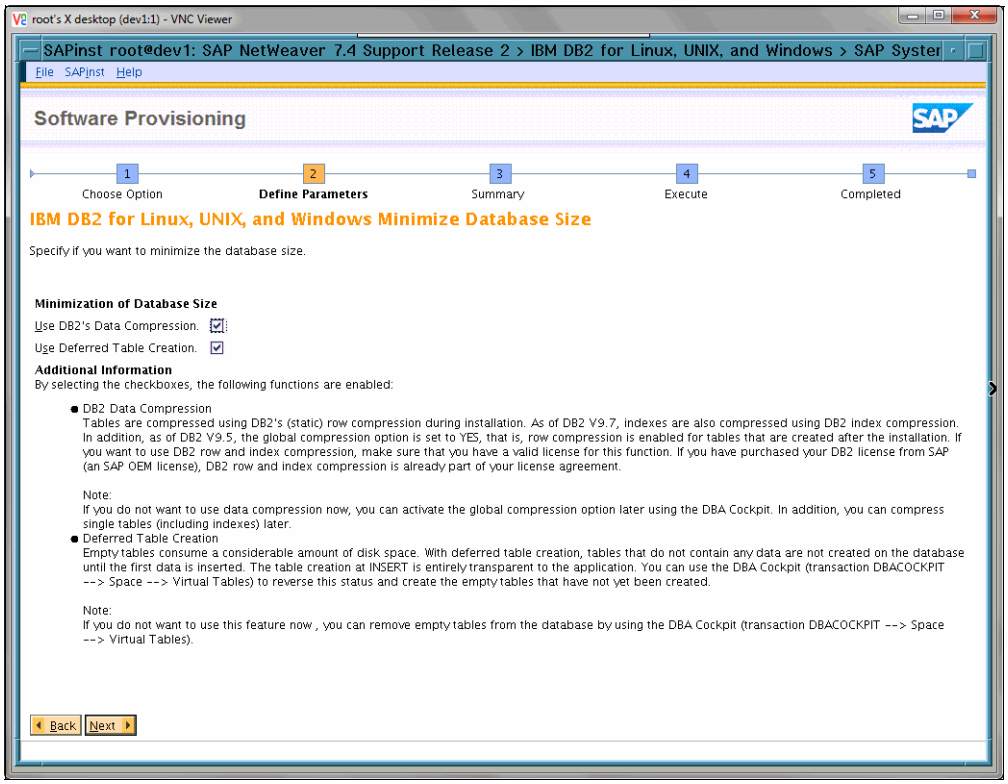

Figure 4-6 Enabling DB2 data compression

Note: The global compression option is ignored due to a regression in dbsl for DB2 for Linux, UNIX, and Windows, when using INTRA_PARALELL=YES. As discussed in 4.2.4, “Enabling intrapartition parallelism for BLU” on page 51, this parameter is necessary to enable BLU. To fix this problem, patch your dbsl as described in the SAP Note 2066163 (DB6: Global Compression Option Set Without Effect).

To get an overview on the compression state of your tables, you can use the catalog view *syscat.tables* and select columns *COMPRESSION*, *TYPE,* and *ROWCOMPMODE* as shown in Example 4-2.

Example 4-2 SAP installation with enabled data compression

```
db2 "select COMPRESSION,ROWCOMPMODE,TYPE,count(*) as COUNT from SYSCAT.TABLES
group by COMPRESSION,ROWCOMPMODE,TYPE"

COMPRESSION ROWCOMPMODE TYPE COUNT
----------- ----------- ---- -----------
                        A              1
B           A           T           6526
```

```
N                        T            362
N                        V          19414
R            A           T            460
V                        T             14
```

Table 4-1 shows a description of the column values.[4]

Table 4-1 Catalog view of the syscat.tables

Type of table (TYPE)	Row compression mode (ROWCOMPMODE)	Compression state (COMPRESSION)
A=Alias	A=Adaptive	B= Both
T=Table (untyped)	S=Static	N=No
V=View	Blank=Row Compression is not enabled.	R= Row Compression
		V= Value Compression
		Blank= Not applicable

In Example 4-2 on page 43, you can see that nearly 7000 tables are already in adaptive row compress mode.

The index compression status can be checked with the command as shown in Example 4-3.

Example 4-3 SAP installation with data compression enabled for the indexes

```
db2 "select COMPRESSION,INDEXTYPE,count(*) as COUNT from SYSCAT.INDEXES group by
COMPRESSION,INDEXTYPE"

COMPRESSION INDEXTYPE COUNT
----------- --------- -----------
N           CLUS               17
N           DIM                23
N           REG              1078
Y           REG              8581
N           XPTH                1
N           XRGN                1
```

Check the SYSCAT.INDEXES catalog view as shown in Table 4-2 for more detailed information.[5]

Table 4-2 Checking the SYSCAT.INDEXES catalog view

Type of index (INDEXTYPE)	Compression state (COMPRESSION)
CLUS=Clustering index	Y=Activated
DIM=Dimension block index	N=Not activated
REG=Regular index	

[4] See http://ibm.co/1OBOCFw
[5] See http://ibm.co/19UKSh3

4.2.2 Enabling automatic storage management during the installation

As a prerequisite for the BLU feature, it is mandatory to enable the automatic storage management (ASM) option. This can be activated during the system installation or SAP import, as shown in Figure 4-7.

Note: The following screen capture is an example of the integration between IBM DB2 BLU and SAP BW.

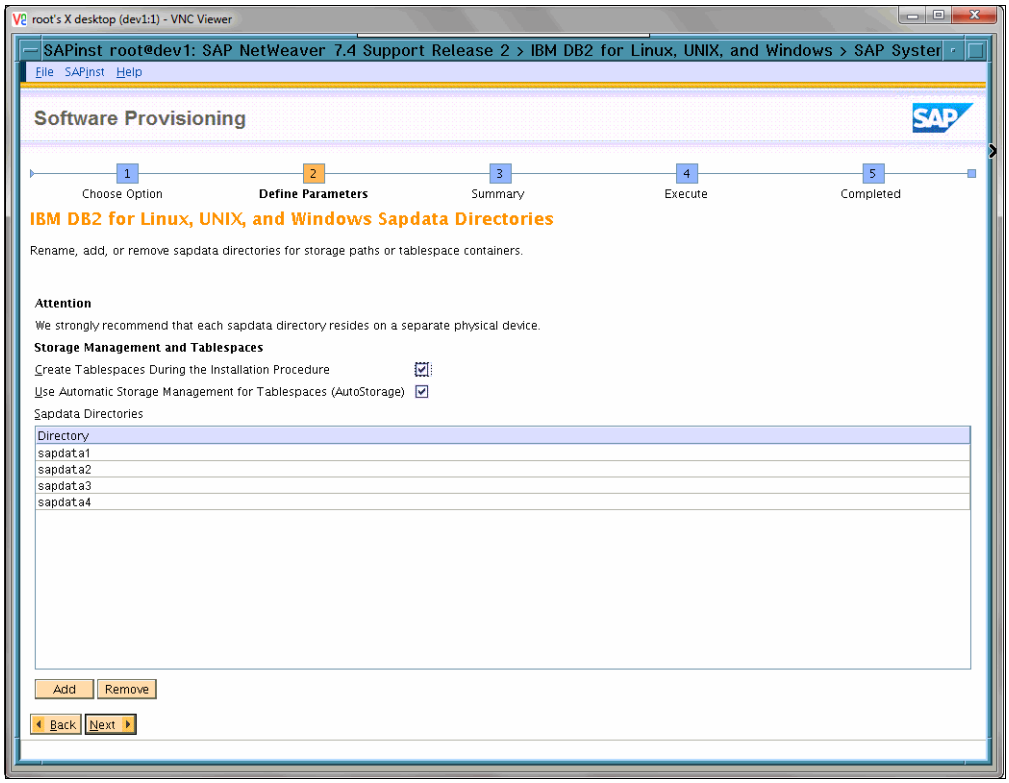

Figure 4-7 Enabling AutoStorage

You can check if ASM is enabled with the command as shown in the Example 4-4.

Example 4-4 SAP BW DEV installation with AutoStorage enabled

```
db2pd -db DEV -storagepaths

Database Member 0 -- Database DEV -- Active -- Up 5 days 16:52:44 -- Date
2014-10-29-09.59.13.262576

Storage Group Configuration:
Address            SGID  Default  DataTag    Name
0x0A00020011165820 0     Yes      0          IBMSTOGROUP

Storage Group Statistics:
Address            SGID  State      Numpaths  NumDropPen
0x0A00020011165820 0     0x00000000 4         0

Storage Group Paths:
Address            SGID  PathID    PathState    PathName
```

```
0x0A00020011189000 0     0         InUse        /db2/DEV/sapdata1
0x0A0002001118B000 0     1         InUse        /db2/DEV/sapdata2
0x0A0002001118D000 0     2         InUse        /db2/DEV/sapdata3
0x0A0002001118F000 0     3         InUse        /db2/DEV/sapdata4
```

For an already installed system without ASM, the activation is explained in the following sections.

4.2.3 Enabling ASM after the SAP BW installation

If you miss the activation of ASM during the installation or your existing SAP BW System, proceed as follows. Remember that without ASM, you cannot use the DB2 BLU feature. The following figures and examples show how to activate ASM for an existing running SAP BW Environment. [6]

Checking the storage paths for ASM

To check if ASM is available or not, you can also use the command as shown in Example 4-5.

Example 4-5 Checking the number of storage paths in the SAP BW System PRD

```
db2 get snapshot for database on PRD |grep "automatic storage"
Number of automatic storage paths          = 0
```

Activating the storage paths for ASM

For the SAP BW installation PRD, ASM was not activated. To activate ASM for the existing storage paths sapdata1-4, proceed as shown in Example 4-6.

Example 4-6 Enabling ASM with an alter database

```
db2 "alter database PRD add storage on '/db2/PRD/sapdata1',  '/db2/PRD/sapdata2',
'/db2/PRD/sapdata3',  '/db2/PRD/sapdata4'"
DB20000I  The SQL command completed successfully.
```

You can check the settings again with the following two commands and verify the change to the number of storage paths for PRD, as shown in Example 4-7 and in Example 4-8 on page 47.

Example 4-7 Storage paths after enabling ASM for PRD

```
db2pd -db PRD -storagepaths

Database Member 0 -- Database PRD -- Active -- Up 1 days 14:32:04 -- Date
2014-10-29-10.12.16.809530

Storage Group Configuration:
Address            SGID  Default  DataTag    Name
0x0A0002003F201000 0     Yes      0          IBMSTOGROUP

Storage Group Statistics:
Address            SGID  State      Numpaths  NumDropPen
0x0A0002003F201000 0     0x00000000 4         0

Storage Group Paths:
```

[6] See http://bit.ly/1NjBMqd

```
Address            SGID  PathID    PathState    PathName
0x0A0002003F204000 0     0         NotInUse     /db2/PRD/sapdata1
0x0A0002003F206000 0     1         NotInUse     /db2/PRD/sapdata2
0x0A0002003F208000 0     2         NotInUse     /db2/PRD/sapdata3
0x0A0002003F251000 0     3         NotInUse     /db2/PRD/sapdata4
```

Example 4-8 ASM number of storage paths changed

```
db2 get snapshot for database on PRD |grep "automatic storage"
Number of automatic storage paths          = 4
```

Converting all DMS table spaces to use the automatic storage

Each DMS table space has to be converted to make use of the ASM feature. You can do this with the command as shown in Example 4-9.

Example 4-9 Converting the DMS table space to use the ASM feature

```
ALTER TABLESPACE <TablespaceName> MANAGED BY AUTOMATIC STORAGE
```

To speed up the conversion task, generate a script for this to convert DSM to ASM table spaces as shown in Example 4-10.

Example 4-10 Generate the script

```
db2 -x "select 'alter tablespace ' || CHR(34) || TBSP_NAME || CHR(34) || ' managed
by automatic storage;' from SYSIBMADM.SNAPTBSP where TBSP_USING_AUTO_STORAGE != 1
and TBSP_TYPE = 'DMS' and TBSP_CONTENT_TYPE in ('ANY', 'LARGE')" > alt_tbs_ams.sql
```

The output of the script generation is shown in Example 4-11.

Example 4-11 Output script generation

```
alter tablespace "SYSCATSPACE" managed by automatic storage;
alter tablespace "SYSTOOLSPACE" managed by automatic storage;
alter tablespace "PRD#POOLD" managed by automatic storage;
alter tablespace "PRD#POOLI" managed by automatic storage;
alter tablespace "PRD#ES74OD" managed by automatic storage;
alter tablespace "PRD#ES74OI" managed by automatic storage;
alter tablespace "PRD#EL74OD" managed by automatic storage;
alter tablespace "PRD#EL74OI" managed by automatic storage;
alter tablespace "PRD#STABD" managed by automatic storage;
alter tablespace "PRD#STABI" managed by automatic storage;
alter tablespace "PRD#BTABD" managed by automatic storage;
alter tablespace "PRD#BTABI" managed by automatic storage;
alter tablespace "PRD#CLUD" managed by automatic storage;
alter tablespace "PRD#CLUI" managed by automatic storage;
alter tablespace "PRD#DIMD" managed by automatic storage;
alter tablespace "PRD#DIMI" managed by automatic storage;
alter tablespace "PRD#FACTD" managed by automatic storage;
alter tablespace "PRD#FACTI" managed by automatic storage;
alter tablespace "PRD#ODSD" managed by automatic storage;
alter tablespace "PRD#ODSI" managed by automatic storage;
alter tablespace "PRD#DDICD" managed by automatic storage;
alter tablespace "PRD#DDICI" managed by automatic storage;
alter tablespace "PRD#DOCUD" managed by automatic storage;
alter tablespace "PRD#DOCUI" managed by automatic storage;
```

```
alter tablespace "PRD#LOADD" managed by automatic storage;
alter tablespace "PRD#LOADI" managed by automatic storage;
alter tablespace "PRD#PROTD" managed by automatic storage;
alter tablespace "PRD#PROTI" managed by automatic storage;
alter tablespace "PRD#SOURCED" managed by automatic storage;
alter tablespace "PRD#SOURCEI" managed by automatic storage;
alter tablespace "PRD#USER1D" managed by automatic storage;
alter tablespace "PRD#USER1I" managed by automatic storage;
```

The execution can be done as shown in Example 4-12.

Example 4-12 Executing the conversion

```
db2 -tvf alt_tbs_ams.sql > alt_tbs_ams.log
```

After the conversion, DMS and AMS table spaces exist in the system as shown in Figure 4-8.

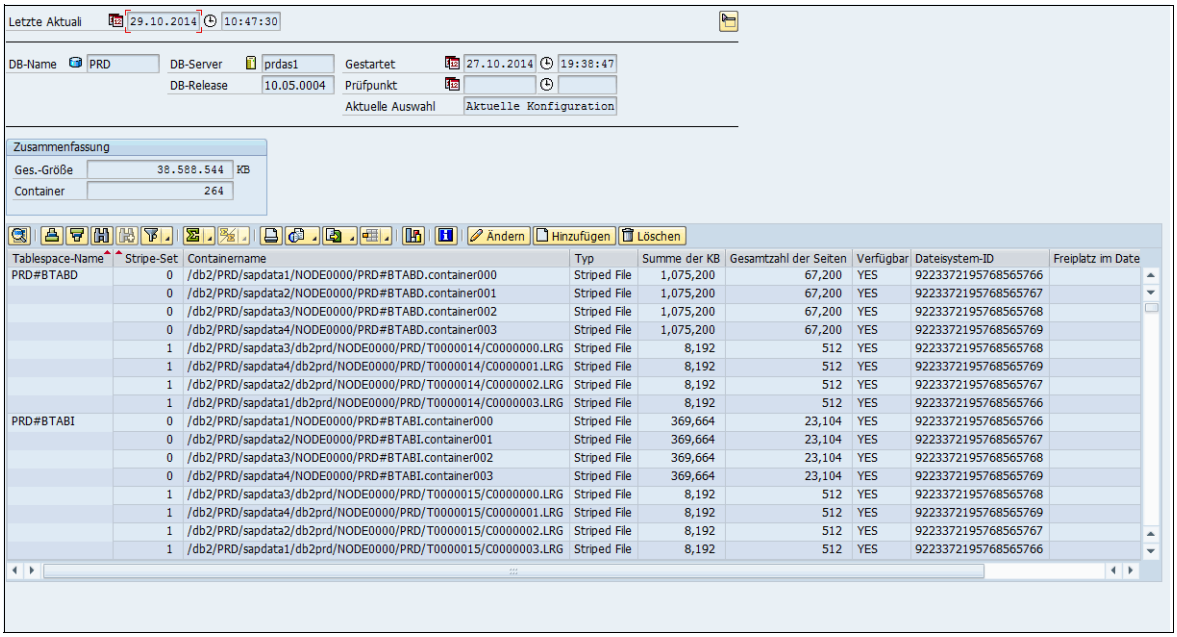

Tablespace-Name	Stripe-Set	Containername	Typ	Summe der KB	Gesamtzahl der Seiten	Verfügbar	Dateisystem-ID	Freiplatz im Date
PRD#BTABD	0	/db2/PRD/sapdata1/NODE0000/PRD#BTABD.container000	Striped File	1,075,200	67,200	YES	9223372195768565766	
	0	/db2/PRD/sapdata2/NODE0000/PRD#BTABD.container001	Striped File	1,075,200	67,200	YES	9223372195768565767	
	0	/db2/PRD/sapdata3/NODE0000/PRD#BTABD.container002	Striped File	1,075,200	67,200	YES	9223372195768565768	
	0	/db2/PRD/sapdata4/NODE0000/PRD#BTABD.container003	Striped File	1,075,200	67,200	YES	9223372195768565769	
	1	/db2/PRD/sapdata3/db2prd/NODE0000/PRD/T0000014/C0000000.LRG	Striped File	8,192	512	YES	9223372195768565768	
	1	/db2/PRD/sapdata4/db2prd/NODE0000/PRD/T0000014/C0000001.LRG	Striped File	8,192	512	YES	9223372195768565769	
	1	/db2/PRD/sapdata2/db2prd/NODE0000/PRD/T0000014/C0000002.LRG	Striped File	8,192	512	YES	9223372195768565767	
	1	/db2/PRD/sapdata1/db2prd/NODE0000/PRD/T0000014/C0000003.LRG	Striped File	8,192	512	YES	9223372195768565766	
PRD#BTABI	0	/db2/PRD/sapdata1/NODE0000/PRD#BTABI.container000	Striped File	369,664	23,104	YES	9223372195768565766	
	0	/db2/PRD/sapdata2/NODE0000/PRD#BTABI.container001	Striped File	369,664	23,104	YES	9223372195768565767	
	0	/db2/PRD/sapdata3/NODE0000/PRD#BTABI.container002	Striped File	369,664	23,104	YES	9223372195768565768	
	0	/db2/PRD/sapdata4/NODE0000/PRD#BTABI.container003	Striped File	369,664	23,104	YES	9223372195768565769	
	1	/db2/PRD/sapdata3/db2prd/NODE0000/PRD/T0000015/C0000000.LRG	Striped File	8,192	512	YES	9223372195768565768	
	1	/db2/PRD/sapdata4/db2prd/NODE0000/PRD/T0000015/C0000001.LRG	Striped File	8,192	512	YES	9223372195768565769	
	1	/db2/PRD/sapdata2/db2prd/NODE0000/PRD/T0000015/C0000002.LRG	Striped File	8,192	512	YES	9223372195768565767	
	1	/db2/PRD/sapdata1/db2prd/NODE0000/PRD/T0000015/C0000003.LRG	Striped File	8,192	512	YES	9223372195768565766	

Figure 4-8 Overview of the table spaces

Converting data to ASM

To get rid of the DMS table spaces, you have to execute a rebalance for each table space as shown in Example 4-13.

Example 4-13 Rebalancing existing table space to AMS

```
ALTER TABLESPACE <TablespaceName> REBALANCE;
```

Example 4-14 shows part of the output for some table spaces.

Example 4-14 Rebalancing output for some table spaces

```
alter tablespace "PRD#POOLD" REBALANCE
DB20000I  The SQL command completed successfully.

alter tablespace "PRD#POOLI" REBALANCE
```

```
DB20000I  The SQL command completed successfully.

alter tablespace "PRD#ES740D" REBALANCE
DB20000I  The SQL command completed successfully.

alter tablespace "PRD#ES740I" REBALANCE
DB20000I  The SQL command completed successfully.
```

You can check the rebalance process by executing the **list** utility command as shown in Example 4-15.

Example 4-15 Checking the rebalance process

```
db2 list utilities show detail

ID                                = 246
Type                              = REBALANCE
Database Name                     = PRD
Member Number                     = 0
Description                       = Tablespace ID: 14
Start Time                        = 10/29/2014 18:11:54.051915
State                             = Executing
Invocation Type                   = User
Throttling:
   Priority                       = Unthrottled
Progress Monitoring:
   Estimated Percentage Complete = 81
      Total Work                  = 37974 extents
      Completed Work              = 30852 extents
      Start Time                  = 10/29/2014 18:11:54.051923

...
```

Temp table spaces in DMS are missing

Figure 4-9 shows that temporary table spaces cannot be converted to ASM with the alter table space.

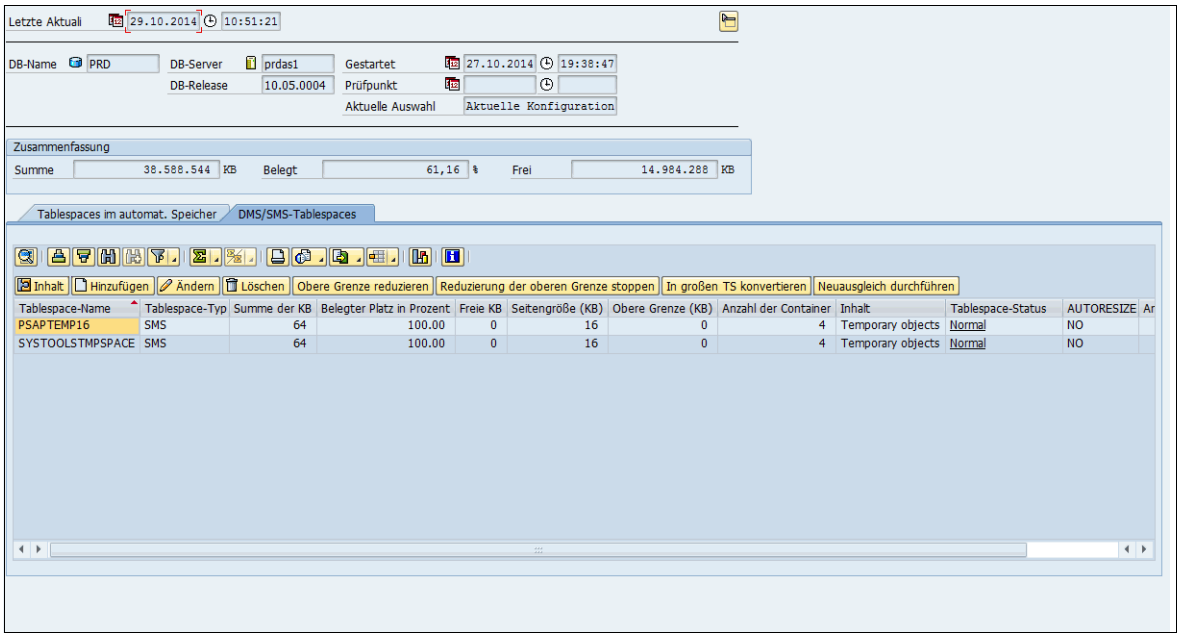

Figure 4-9 Temporary table spaces still in DMS format

To use ASM for temp table space, you have to re-create them. The entire process for PSAPTEMP16 and SYSTOOLTMPSAPCE is shown in Example 4-16.

Example 4-16 Procedure to convert Temp table spaces to ASM

```
db2 “rename tablespace PSAPTEMP16 to oldTEMP16“
DB20000I  The SQL command completed successfully.

db2 “rename tablespace SYSTOOLSTMPSPACE to oldTOOLSTMPSPACE“
DB20000I  The SQL command completed successfully.

db2 “create temporary tablespace PSAPTEMP16 in nodegroup IBMTEMPGROUP pagesize 16k
extentsize 2 prefetchsize automatic no file system caching dropped table recovery
off“
DB20000I  The SQL command completed successfully.

db2 “create user temporary tablespace SYSTOOLSTMPSPACE in nodegroup IBMCATGROUP
pagesize 16k extentsize 2 prefetchsize automatic no file system caching dropped
table recovery off“
DB20000I  The SQL command completed successfully.

db2 “drop tablespace oldTEMP16“
DB20000I  The SQL command completed successfully.

db2 “drop tablespace oldTOOLSTMPSPACE“
DB20000I  The SQL command completed successfully.
```

Figure 4-10 shows that no DMS table spaces are available for the SAP System PRD. For moving temporary table space to multi-temperature storage, see 4.2.7, “Multi-temperature storage groups for DB2 temp table spaces” on page 53.

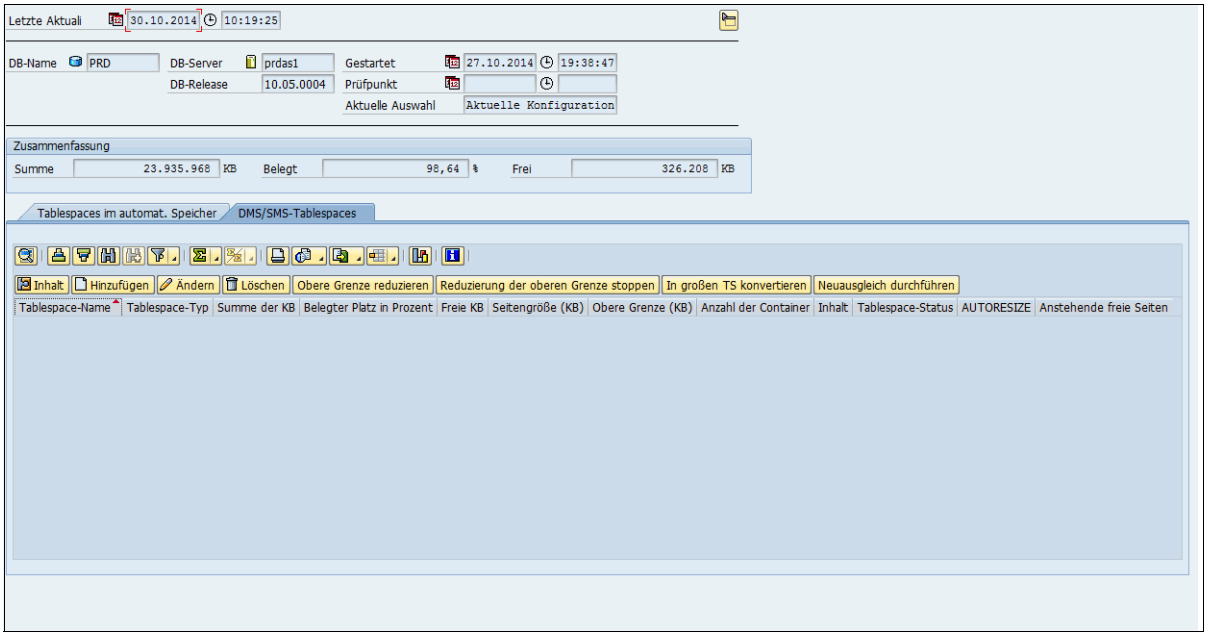

Figure 4-10 No DMS table space in the PRD system

4.2.4 Enabling intrapartition parallelism for BLU

To use BLU Acceleration, enabling intrapartition parallelism is required. Refer to 4.1.4, “DB2 database layout for BLU” on page 41. To do so, you have to change the following database manager parameters:

- INTRA_PARALLEL
- MAX_QUERYDEGREE

As described in SAP Note 1851832, set the following values for both the dbm parameters and restart instance as shown in Example 4-17.[7]

Example 4-17 Enabling intrapartition parallelism for BLU

```
db2 update dbm cfg using INTRA_PARALLEL YES
DB20000I  The UPDATE DATABASE MANAGER CONFIGURATION command completed
successfully.

db2 update dbm cfg using MAX_QUERYDEGREE ANY
DB20000I  The UPDATE DATABASE MANAGER CONFIGURATION command completed
successfully.
```

DB2 BLU is then automatically enabled when the hardware and software requirements are fulfilled. Refer to Chapter 2, “Requirements for an SAP NetWeaver BW environment” on page 17 for more details.

[7] See http://service.sap.com/notes (authentication required)

Note: Documentation recommends setting DB2_WORKLOAD to ANALYTICS to enable BLU Acceleration. However, if you are using BLU Acceleration with SAP BW, all required settings are available in DB2_WORKLOAD=SAP.

4.2.5 DB2 Cancun parameter settings for BLU

When using BLU Acceleration, you have to configure the DB2 parameter SHEAPTHRES_SHR to 40% of the DATABASE_MEMORY and set the SORTHEAP to 1/20 of the SHEAPTHRES_SHR. To assure good BLU performance on production systems, set the INSTANCE_MEMORY to at least 64 GB.

Note: Automatic tuning of sort memory is not supported in environments using column-organized tables (BLU Acceleration) in the initial release. Self-tuning memory manager (STMM) will not be able to recognize the costs and benefits of sort memory usage during column-organized data processing, and might undersize the sort memory configuration as a result. This can lead to potential performance issues or query failures. Do not set SHEAPTHRES_SHR and SORTHEAP to AUTOMATIC when using BLU Acceleration. Other memory consumers are not affected and should be set to AUTOMATIC.

Loading data into BLU tables with the DB2 load utility, for example using R3load, requires sufficient memory in the utility heap to achieve optimal compression. In regular configurations, set UTIL_HEAP_SZ to AUTOMATIC. This allows the utility heap to grow into the overflow area of DATABASE_MEMORY. In a small test or in a QA system configuration with less than 64 GB memory, the AUTOMATIC setting might not be large enough. In such cases, set UTIL_HEAP_SZ to 2 GB before loading the R3load to avoid suboptimal compression.

It is recommended that you enable STMM. However, if you do not want to use STMM, you should set DATABASE_MEMORY to COMPUTED and configure your buffer pools to about the same size as SHEAPTHRES_SHR.

For your SAP BW landscape, the following DB2 parameter configuration is established as shown in Table 4-3.

Table 4-3 BLU DB2 parameter configuration for SAP BW landscape

DB2 parameter	DEV	QAS	PRD
INSTANCE_MEMORY	20 GB	20 GB	48 GB
DATABASE_MEMORY	AUTOMATIC	AUTOMATIC	AUTOMATIC
SHEAPTHRES_SHR	8 GB	8 GB	25,6 GB
SORTHEAP	0,4 GB	0,4 GB	1,28 GB
UTIL_HEAP_SZ	50000 pages	50000 pages	AUTOMATIC

For more details, check the SAP Note 1851832 (DB6: DB2 10.5 Standard Parameter).

Note: Belonging to the NUMA effect, described in 3.5, “Active memory expansion is not recommended in this case” on page 34, currently there is no official statement from SAP whether to enable the DB2 affinitization support through setting the DB2 registry variable DB2_RESOURCE_POLICY to AUTOMATIC.

4.2.6 SAP BW RSADMIN parameters for DB2 BLU Acceleration

For new SAP BW installations, set the following RSADMIN parameter as shown in Table 4-4 to create the SAP BW temporary tables, PSA tables, and InfoObject tables as column-organized tables.

Table 4-4 SAP BW RSADMIN parameters for DB2 BLU Acceleration

RSADMIN parameter name and value	Comment
DB6_IOBJ_USE_CDE=YES	Create SID and time-independent and time-dependent attribute SID tables of new characteristic InfoObjects as column-organized tables.
DB6_PSA_USE_CDE=YES	Create new PSA tables as column-organized tables.
DB6_TMP_USE_CDE=YES	Create new BW temporary tables as column-organized tables.

Note: SAP Note 2020190 shows that the performance issues are solved when a customer is using the RSADMIN parameter **`mass_cube_writer`**.

For more information, refer to the *Architecting and Developing DB2 with BLU Acceleration*, SG24-8212:

http://www.redbooks.ibm.com/abstracts/sg248212.html

4.2.7 Multi-temperature storage groups for DB2 temp table spaces

When activating ASM during or after the installation, as described in the previous sections, only one storage path is defined. This one group is used by all ASM table spaces that are defined in the database.[8]

As a consequence of the trend to big data in the business world, more storage capacity and different storage types are required. With the use of multi-temperature storage, you can add more storage types via different storage paths.

Generally, the access to big data has to be fast. To ensure this, it is necessary to distribute the business data on different storage media based on their access time or temperature, see 1.6, “Multi-temperature storage group” on page 13 for a detailed description.

One common use for multi-temperature storage in an SAP BW environment is to divide the data and temporary table spaces in different storage groups. This is also highly recommended by SAP.

[8] See http://bit.ly/1FMpoju

In the next section, the following example shows how to move temporary table spaces to new storage paths.

Prerequisites

As described in 1.6, “Multi-temperature storage group” on page 13, multi-temperature storage is supported since DB2 for Linux, UNIX, and Windows Version 10.1. The DB2 database furthermore has to be enabled for automatic storage as described before in 4.3.2, “Backup and restore” on page 57.

Checking and adding a storage group

First, check which storage groups and paths are available. Example 4-18 shows that one storage group IBMSTOGROUP is available, and all the storage paths from sapdata1-4 are assigned to it.

Example 4-18 Storage group and storage paths

```
db2 "SELECT VARCHAR(storage_group_name, 30) AS STOGROUP,
VARCHAR(db_storage_path,40) AS STORAGE_PATH FROM
TABLE(ADMIN_GET_STORAGE_PATHS('',-2)) AS T"

STOGROUP                       STORAGE_PATH
------------------------------ ----------------------------------------
IBMSTOGROUP                    /db2/DEV/sapdata1
IBMSTOGROUP                    /db2/DEV/sapdata2
IBMSTOGROUP                    /db2/DEV/sapdata3
IBMSTOGROUP                    /db2/DEV/sapdata4
```

To add a storage group, use the command as shown in Example 4-19.

Example 4-19 Adding a storage group

```
db2 "CREATE STOGROUP TMPSTOGROUP ON
'/db2/DEV/sapdata5','/db2/DEV/sapdata6','/db2/DEV/sapdata7','/db2/DEV/sapdata8'"
DB20000I  The SQL command completed successfully.
```

The file system sapdata5-8 can be on another storage type. Example 4-20 shows the addition of internal SSD flash storage.

Example 4-20 Adding an internal SSD flash storage

```
lscfg -l hdisk4
  hdisk4           U8233.E8B.100C99R-V4-C4-T1-L8200000000000000  Virtual SCSI
Solid State Drive
```

Example 4-21 verifies that two storage groups are available now.

Example 4-21 Check storage groups

```
db2 "SELECT sgname FROM syscat.stogroups"

SGNAME
-------------------
IBMSTOGROUP
TMPSTOGROUP
```

Moving temporary table spaces to a new storage group

To move temporary table spaces to a new storage group, define a new temp table space PSAPTEMP_16 and remove the table from prior storage groups.

To create a new temporary table space PSAPTEMP, proceed as shown in Example 4-22.

Example 4-22 New temporary table space in PSAPTEMP_16

```
db2 "create temporary tablespace PSAPTEMP_16 in nodegroup IBMTEMPGROUP pagesize
16k MANAGED BY AUTOMATIC STORAGE extentsize 2 prefetchsize automatic no file
system caching dropped table recovery off using stogroup tmpstogroup"
DB20000I  The SQL command completed successfully.
```

Then, drop the temporary table space PSAPTEMP16 in the IBMSTOGROUP, as shown in Example 4-23.

Example 4-23 Dropping table space PSAPTEMP16 in IBMSTOGROUP

```
db2 drop tablespace PSAPTEMP16
DB20000I  The SQL command completed successfully.
```

Then, create PSAPTEMP16 in the new storage group TMPSTOGROUP as shown in Example 4-24, and drop PSAPTEMP_16.

Example 4-24 Creating the table space

```
db2 "create temporary tablespace PSAPTEMP16 in nodegroup IBMTEMPGROUP pagesize 16k
MANAGED BY AUTOMATIC STORAGE extentsize 2 prefetchsize automatic no file system
caching dropped table recovery off using stogroup tmpstogroup"
DB20000I  The SQL command completed successfully.

db2 drop tablespace PSAPTEMP_16
DB20000I  The SQL command completed successfully.
```

The same procedure can be done for the SYSTOOLTMPSPACE. Example 4-25 shows the possibility to rename the old temporary table space SYSTOOLTMPSPACE before creating a table space and dropping the old one.

Example 4-25 Procedure for SYSTOOLTMPSPACE

```
db2 rename tablespace SYSTOOLSTMPSPACE to oldTOOLSTMPSPACE
DB20000I  The SQL command completed successfully.

db2 create user temporary tablespace SYSTOOLSTMPSPACE in nodegroup IBMCATGROUP
pagesize 16k extentsize 2 prefetchsize automatic no file system caching dropped
table recovery off
DB20000I  The SQL command completed successfully.

db2 drop tablespace oldTOOLSTMPSPACE
DB20000I  The SQL command completed successfully.
```

Note: For performance reasons, create the same storage paths for data and temporary table spaces, as shown in Example 4-26 on page 56.

Example 4-26 Creating storage paths for data and temporary table spaces

```
db2 "SELECT VARCHAR(storage_group_name, 30) AS STOGROUP,
VARCHAR(db_storage_path,40) AS STORAGE_PATH FROM
TABLE(ADMIN_GET_STORAGE_PATHS('',-2)) AS T"

STOGROUP                       STORAGE_PATH
------------------------------ ----------------------------------------
IBMSTOGROUP                    /db2/DEV/sapdata1
IBMSTOGROUP                    /db2/DEV/sapdata2
IBMSTOGROUP                    /db2/DEV/sapdata3
IBMSTOGROUP                    /db2/DEV/sapdata4
TMPSTOGROUP                    /db2/DEV/sapdata5
TMPSTOGROUP                    /db2/DEV/sapdata6
TMPSTOGROUP                    /db2/DEV/sapdata7
TMPSTOGROUP                    /db2/DEV/sapdata8
```

4.3 Administration and operation with SAP DB2 BLU

This section shows how transparent DB2 BLU fits into the system landscape and housekeeping remains unchanged.

4.3.1 Transport system

If all the systems in the SAP BW landscape fulfill the BLU requirements, *R3trans* and *tp* recognize automatically whether the target BLU object is row or column-organized. For example, if the source BLU object is column-organized and the target BLU object is row-organized, the object is transported as a row-organized table in the target. Figure 4-11 shows a typical three-system landscape.

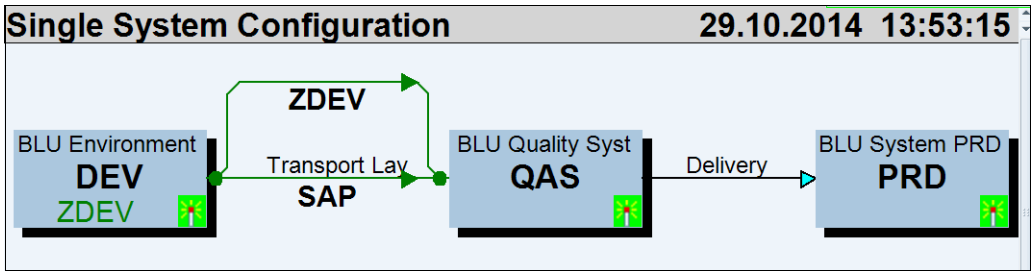

Figure 4-11 SAP standard three-system landscape

Check that all mandatory prerequisites for BLU mentioned in SAP Note 2019648 (DB6: Transport of a Column-Organized InfoCube Fails) are available in all systems belonging to the transport system to avoid transport errors:

- BLU-supported DB2 version is in all systems.
- ASM is activated.
- Reclaimable storage as well.

4.3.2 Backup and restore

BLU is deeply integrated in DB2 10.5. So there are no changes to simple backups or restores. Nevertheless, you should optimize your backup strategy for big data systems. In the next section, you find some considerations on backups and restores:

- The multi-temperature strategy can provide benefits for backup performance and reduce the backup volume and backup workload on SAP Business Warehouse systems. For example, using the Near-Line Storage feature for cold temperature data moves the rarely used data to an online archive system. Refer to Chapter 5, “Adding IBM DB2 Near-Line Storage to archive cold temperature SAP BW data” on page 59.
- The DB2 backup utility supports full or incremental database backups. To benefit from higher parallelism on the table space level, distribute your data across smaller multiple table spaces. To do so, you can use Report DB6CONV for ABAP Systems to convert tables to new table spaces.
- Incremental backups help to reduce backup volume. This can be enabled with the database configuration parameter **TRACKMOD=YES**. DB2 then tracks only changed pages that must be included in an incremental backup. Table spaces that did not change since the last incremental backup are entirely skipped during the backup processing and there is no need to check on the page level, which results in faster backup execution.

4.3.3 System copies

Several scenarios are described in the SAP Note 886102 (System Landscape Copy for SAP NetWeaver BW) about how to establish system copies in an SAP BW landscape. You can use the redirected restore for system copies without any additional configuration. If you want to use R3load, you have also the possibility to migrate directly to BLU tables with the use of Report SMIGR_CREATE_DDL.

For more information about BLU and SAP Migration, refer to section 8.7.3 of *Architecting and Developing DB2 with BLU Acceleration*, SG24-8212:

http://www.redbooks.ibm.com/abstracts/sg248212.html

5

Adding IBM DB2 Near-Line Storage to archive cold temperature SAP BW data

As described in 1.6, “Multi-temperature storage group” on page 13, not all data is used with the same frequency. The amount of data found in analytic systems such as an SAP Business Warehouse is growing constantly, but especially the older data, which is used rarely and can be assigned to cold temperature data. With the Near-Line Storage (NLS) feature, you can archive older but needed data into a separate database on another LPAR and also to other storage subsystems. You can access the data online and also accelerated with DB2 BLU for SAP Business Warehouse InfoProviders.

This chapter describes different NLS scenarios and considerations for the hardware design. Furthermore, you get a short overview on how to install NLS and connect it to an SAP BW System.

The following topics are discussed in this chapter:

- NLS scenario for SAP BW landscape
- Implementing the Near-Line Storage feature

© Copyright IBM Corp. 2015. All rights reserved.

5.1 NLS scenario for SAP BW landscape

NLS is a category of data persistency that is similar to archiving. With NLS, you can transfer historical read-only data of your InfoProviders to an additional NLS database. For more details, refer to 1.5, “Archiving with NLS” on page 11.

Adding additional NLS databases to the current system landscape can result in a different design approach. The section introduces two main possibilities.[1]

You can add an additional NLS database per each SAP BW System. This additional NLS database can be on the same LPAR with different volume groups and storage system or even on an additional LPAR. Figure 5-1 shows a possible case with an NLS database for each BW system in the landscape.

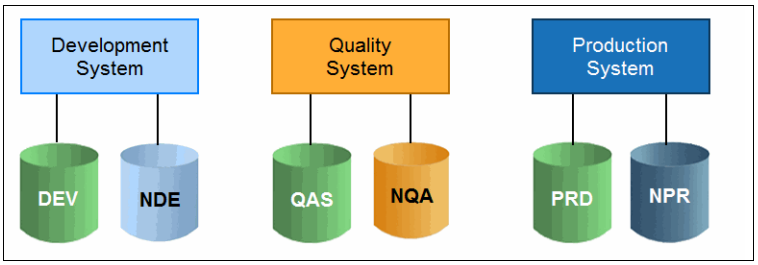

Figure 5-1 NLS database for each SAP BW System

This system configuration has good NLS performance at the database level, and the database recovery is independent, but this results in more administration effort.

To reduce the administration effort, one NLS database for several SAP BW Systems can be a good approach as shown in Figure 5-2.

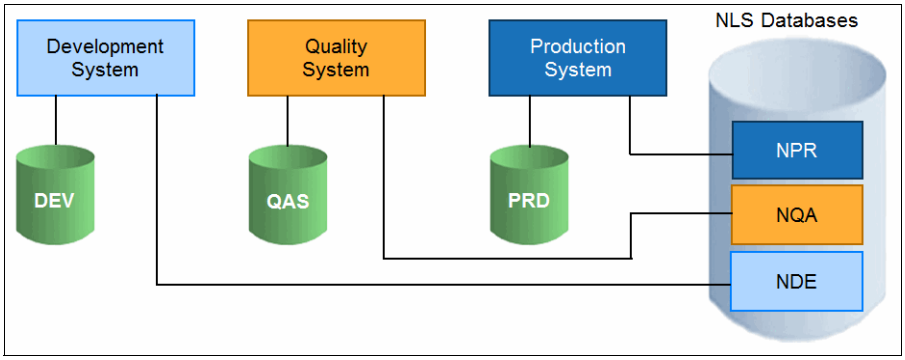

Figure 5-2 One NLS database for several SAP BW Systems

But keep in mind that the performance for the NLS databases can be impacted by other SAP BW applications. When recovering one NLS database, there is no access to the other databases.

[1] See http://www.informatik.uni-jena.de/dbis/lehre/ss2011/dbarch/SAP_BW_DB2_NLS_22062011.pdf

5.1.1 Design for SAP NLS

To have a meaningful NLS database setup, it is recommended to install the production NLS database separately from the QAS/DEV NLS databases. For the preproduction database system, a shared approach for test and quality is reasonable. This scenario takes also into account that preproduction and production systems can be in different data centers. Figure 5-3 shows this approach.

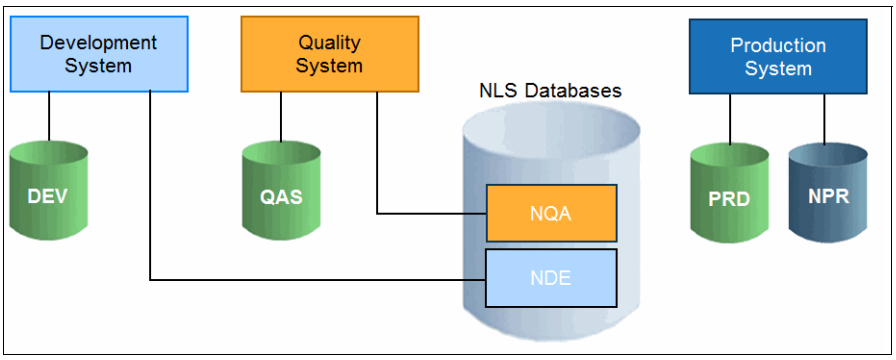

Figure 5-3 One NLS database for preproduction and one for production

The following NLS scenario is based on the approach as shown in Figure 5-3.

Note: Your production system with the DB2 for Linux, UNIX, and Windows database should have at a minimum a compressed size of 5 TB before considering on using NLS.

5.1.2 Minimum resources for SAP NLS

This landscape scenario is sized with the minimum requirements that are recommended in *SAP Note 1851832 (DB6 DB2 10.5 Standard Parameter Settings)*. For NLS, the same minimum size for test and QA system is recommended. This means at least 8 CPU and 32 GB Memory for NDE, NQA, and NPR as shown in Figure 5-4.

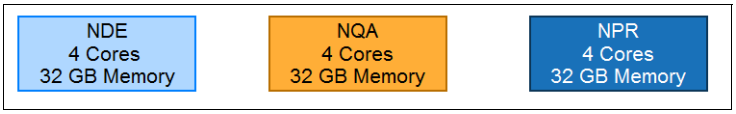

Figure 5-4 CPU and memory baseline for SAP NLS databases

The SAP systems are installed in an IBM POWER7 processor-based server with AIX as described in 4.1, “Design SAP BW landscape” on page 38.

5.1.3 Considerations for the preproduction systems (NDE and NQA)

The development system is usually a small configuration as mentioned in 4.1.2, “Considerations for the preproduction systems (DEV and QAS)” on page 39. For NDE and NQA, the performance is not the main driving factor so that the cold temperature data of NDE and NQA can be moved to an additional LPAR in the same data center. Because the BLU feature is also available on the NLS systems NDE and NQA, each should use four cores and 8 GB memory per core. The CPU resources should be in capped mode so that no further CPU resources are used by the NLS systems as shown in Figure 5-5.

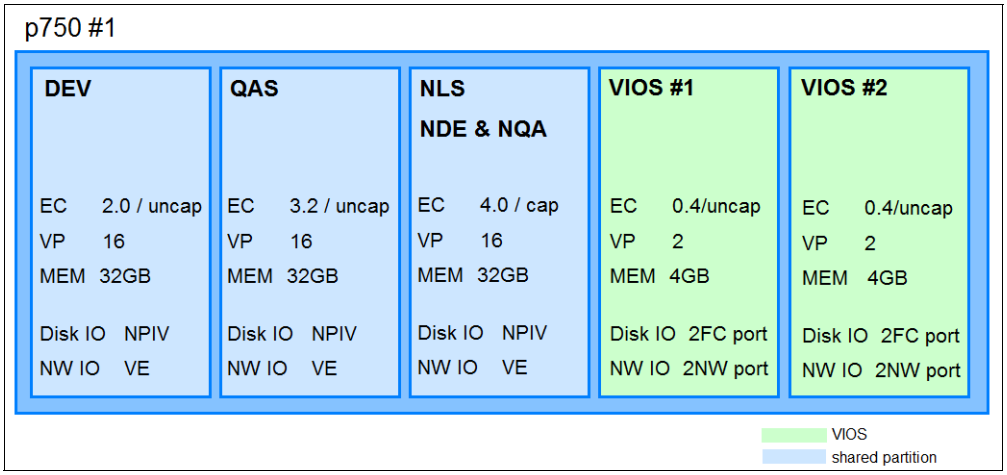

Figure 5-5 LPAR for NLS preproduction systems

5.1.4 Consideration for the production system (NPR)

The NLS production system NPR, first of all, has to deliver a large storage system so that the user can archive big data to the online archive. To guarantee enough BLU performance on NLS and good usage of the CPU caches, there is no need to install NPR in a dedicated POWER environment. In this scenario, we installed NPR on an additional LPAR with capped CPU resources. This approach provides the possibility to run the SAP and NLS environment still in a flexible shared POWER environment as shown in Figure 5-6.

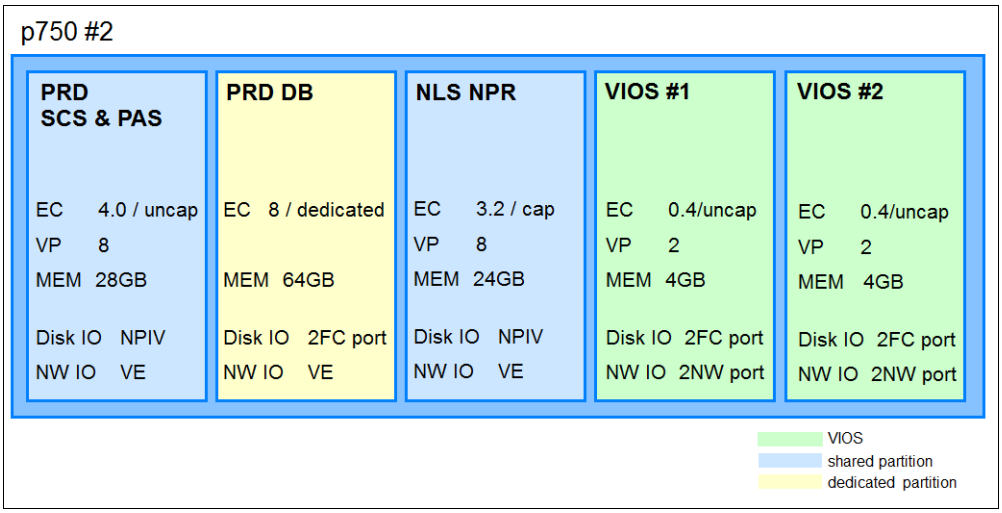

Figure 5-6 LPAR for NLS production system

5.2 Implementing the Near-Line Storage feature

For the entire installation process of NLS databases, refer to the *SAP Guide: Enabling SAP NetWeaver Business Warehouse Systems to Use IBM DB2 for Linux, UNIX, and Windows as Near-Line Storage* (*NLS*), at the following web page:

http://bit.ly/1FRYa7H

The following section shows how to install NLS and how to set up the NLS connection to the SAP BW Systems.

5.2.1 Installation of the NLS database NDE

The SAP Software Provisioning Manager entry point for an NLS database installation can be found under **Generic Installation Options** → **IBM DB2 for Linux, UNIX, and Windows** → **Database Tools** → **Install Near-Line Storage Database**, as shown in Figure 5-7.

Note: The following screen capture is an example of the integration between IBM DB2 BLU and SAP BW.

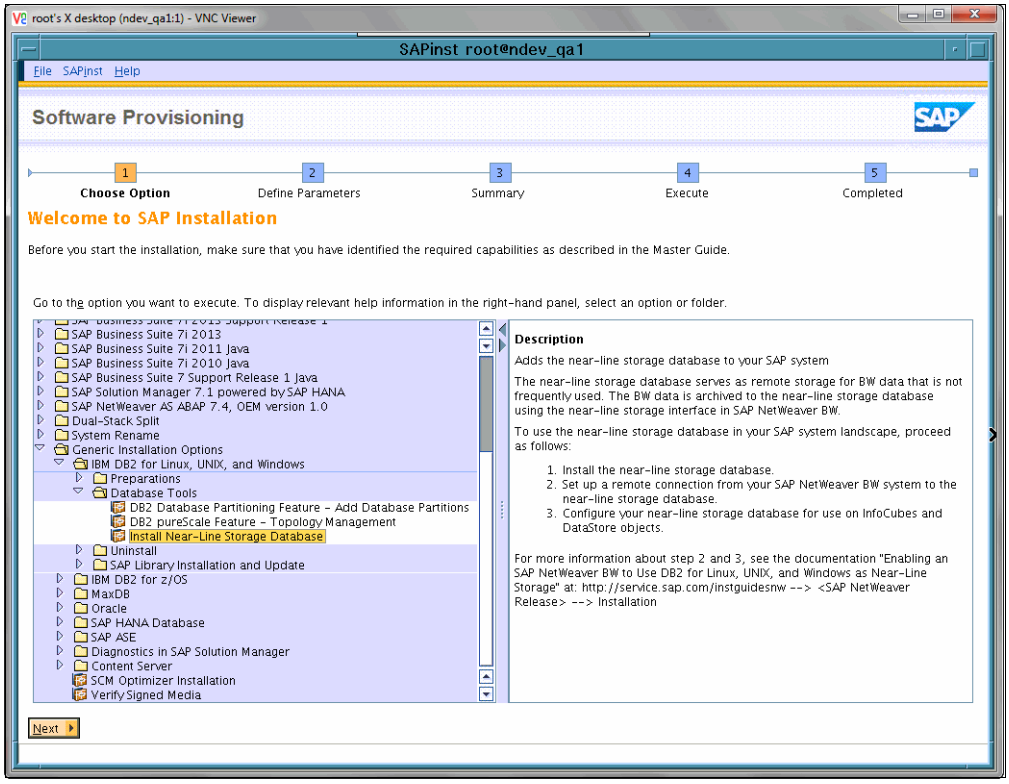

Figure 5-7 SAPinst Entry Point for the NLS database installation

The SAPinst performs the following steps:

- Installs the DB2 software on the appropriate server.
- Creates the operating system users and groups.
- Creates a DB2 instance.
- Creates and activates the DB2 database for NLS.

Table 5-1 lists the important input parameters for NDE, NQA, and NPR, which you must specify during the dialog phase of the SAPinst NLS database installation.

Table 5-1 Input for NLS installation for each system

NDE	NQA	NPR
DBSID: NDE	DBSID: NQA	DBSID: NPR
RDBMS DVD: DB2 V10.5	RDBMS DVD: DB2 V10.5	RDBMS DVD: DB2 V10.5
PW of DB-User db2nde	PW of DB-User db2nqa	PW of DB-User db2npr
OS Groups ID number	OS Groups ID number	OS Groups ID number
DB2 SW Destination: /db2/NDE/db2_software	DB2 SW Destination: /db2/NQA/db2_software	DB2 SW Destination: /db2/NPR/db2_software
DB Communication Port 5912	DB Communication Port 5922	DB Communication Port 5912
Communication between Partitions 5914 - 5917	Communication between Partitions 5924 - 5927	Communication between Partitions 5914 - 5917
DB Instance Memory 10 GB	DB Instance Memory 10 GB	DB Instance Memory 30 GB
Activate ASM (AutoStorage)	Activate ASM (AutoStorage)	Activate ASM (AutoStorage)
Create TSP sapdata1-4	Create TSP sapdata1-4	Create TSP sapdata1-4

Be careful when choosing the database ports for NQA because NDE uses already standard ports on the same LPAR.

For the Near-Line Storage database, you need the following DB2 file systems as shown in Example 5-1.

Example 5-1 NLS file systems for NDE and NQA

```
# df -m
Filesystem    MB blocks      Free %Used    Iused %Iused Mounted on
/dev/hd4         256.00     54.10   79%    10177    43% /
/dev/hd2        2176.00    166.36   93%    41534    50% /usr
/dev/hd9var      448.00    147.53   68%     6331    16% /var
/dev/hd3        2176.00   2170.58    1%       84     1% /tmp
/dev/hd1          64.00     63.65    1%        5     1% /home
/dev/hd11admin    128.00    127.63    1%        5     1% /admin
/proc                -         -    -         -     -  /proc
/dev/hd10opt     320.00    178.96   45%     6996    15% /opt
/dev/livedump    256.00    255.64    1%        4     1% /var/adm/ras/livedump
/dev/download_lv  18496.00   7940.97   58%    13798     1% /download
/dev/work_lv    3200.00   3199.12    1%       11     1% /work
/dev/fslv00     2048.00   1163.35   44%     1825     1% /db2/db2nde
/dev/fslv01     3072.00    576.64   82%     6633     5% /db2/NDE
/dev/fslv04     5120.00   1279.31   76%       67     1% /db2/NDE/log_dir
/dev/fslv05      256.00    247.91    4%       12     1% /db2/NDE/db2dump
/dev/fslv06    13312.00  13261.54    1%       24     1% /db2/NDE/sapdata1
/dev/fslv07    13312.00  13261.54    1%       24     1% /db2/NDE/sapdata2
/dev/fslv08    13312.00  13261.54    1%       24     1% /db2/NDE/sapdata3
/dev/fslv09    13312.00  13261.54    1%       24     1% /db2/NDE/sapdata4
/dev/fslv10     2048.00   1163.33   44%     1822     1% /db2/db2nqa
/dev/fslv11     3072.00    576.64   82%     6633     5% /db2/NQA
/dev/fslv14     5120.00   1279.31   76%       67     1% /db2/NQA/log_dir
```

```
/dev/fslv15      256.00    248.08    4%       12     1% /db2/NQA/db2dump
/dev/fslv16    13312.00  13261.54    1%       24     1% /db2/NQA/sapdata1
/dev/fslv17    13312.00  13261.54    1%       24     1% /db2/NQA/sapdata2
/dev/fslv18    13312.00  13261.54    1%       24     1% /db2/NQA/sapdata3
/dev/fslv19    13312.00  13261.54    1%       24     1% /db2/NQA/sapdata4
```

The file system relies on the IBM DS8000 enterprise storage solution, which provides enough file system space for the databases to grow. Comparing the NLS installation with the SAP BW Installation described in 4.2, “Implementation of SAP BW with DB2 BLU” on page 42, the NLS database is per default an ASM Installation as shown in Example 5-2.

Example 5-2 ASM is enabled per default during the SAP installation procedure

```
db2 get snapshot for database on NDE | grep "automatic storage"
Number of automatic storage paths          = 4
ndev_qa1:db2nde 12>

ndev_qa1:db2nde 1> db2pd -db NDE -storagepaths

Database Member 0 -- Database NDE -- Active -- Up 0 days 00:07:47 -- Date
2014-11-13-05.06.31.335655

Storage Group Configuration:
Address            SGID  Default  DataTag    Name
0x0A00020011165820 0     Yes      0          IBMSTOGROUP

Storage Group Statistics:
Address            SGID  State      Numpaths  NumDropPen
0x0A00020011165820 0     0x00000000 4         0

Storage Group Paths:
Address            SGID  PathID    PathState    PathName
0x0A00020011189000 0     0         InUse        /db2/NDE/sapdata1
0x0A0002001118B000 0     1         InUse        /db2/NDE/sapdata2
0x0A0002001118D000 0     2         InUse        /db2/NDE/sapdata3
0x0A0002001118F000 0     3         InUse        /db2/NDE/sapdata4

db2 get snapshot for database on NQA |grep "automatic storage"
Number of automatic storage paths          = 4
```

But for good performance, you should also divide the data and temp table spaces as described in 4.2.7, “Multi-temperature storage groups for DB2 temp table spaces” on page 53.

5.2.2 Configure the new database connection in the DBACockpit

To establish the NLS database connection to the SAP BW System, you first have to set up a connection in the transaction DBACOCKPIT as shown in Figure 5-8.

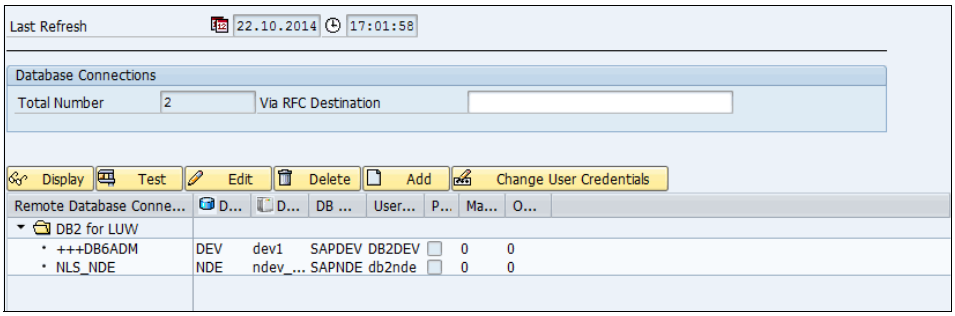

Figure 5-8 NLS_NDE connection for DEV

In this case, we use the connection name NLS_NDE. The connection details are shown in Figure 5-9.

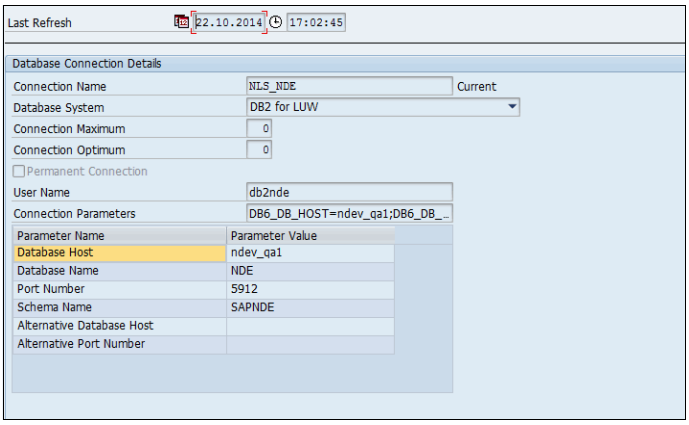

Figure 5-9 Details of NLS_NDE connections

5.2.3 Configure the NLS connection in SPRO

To set up the Near-Line Storage connection for DEV via the NLS_NDE database connection, go to the Transaction SPRO and to **SAP Customizing Implementation Guide** → **SAP NetWeaver** → **Business Warehouse** → **General Settings** → **Process Near-Line Storage Connection**, as shown in Figure 5-10.

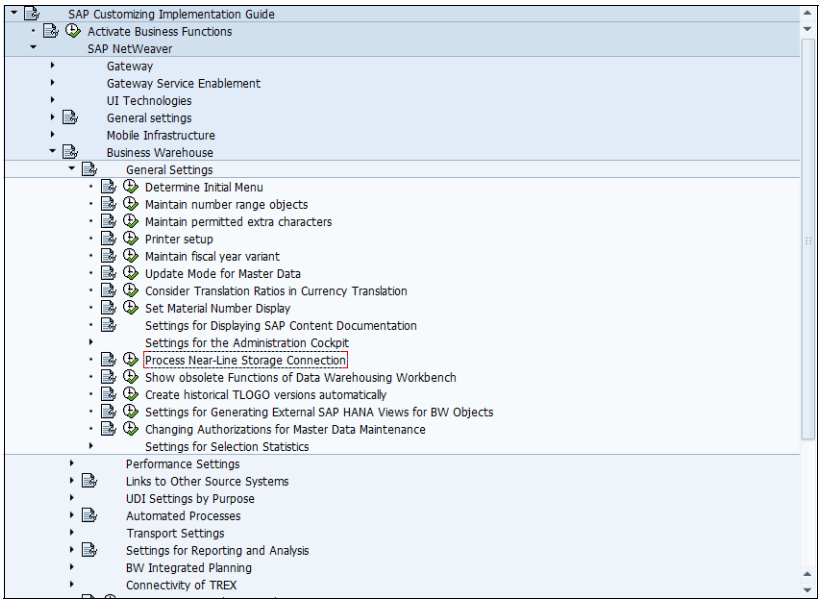

Figure 5-10 Process the NLS connection for SAP BW System DEV

To complete this process for BW archiving, add two parameters as shown in Figure 5-11.

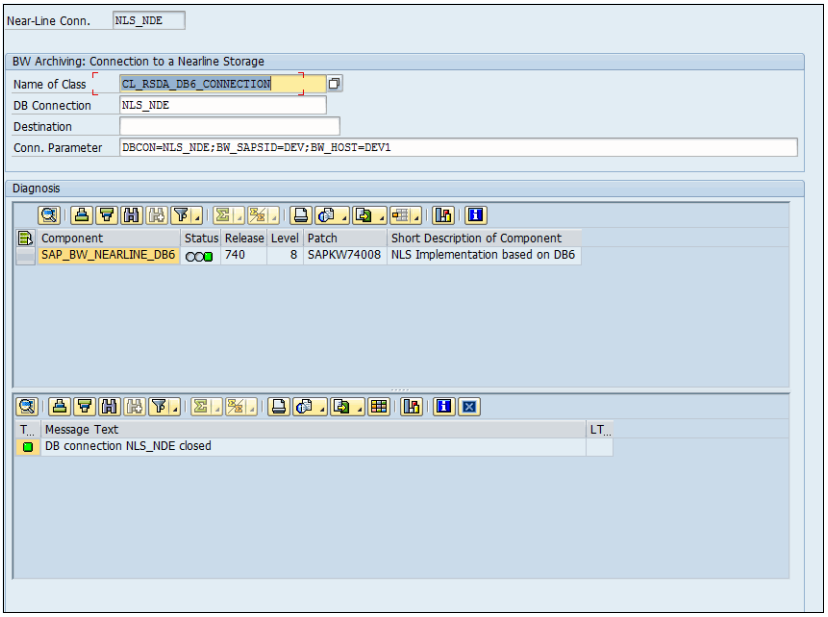

Figure 5-11 Input parameters for the SAP BW archiving connection

With the input values *CL_RSDA_DB6_CONNECTION* and *NLS_NDE*, the *Connection. Parameter* field is generated automatically after executing the transaction. With this setup, you can administer the NLS database NDE in Transaction DB02, as shown in Figure 5-12.

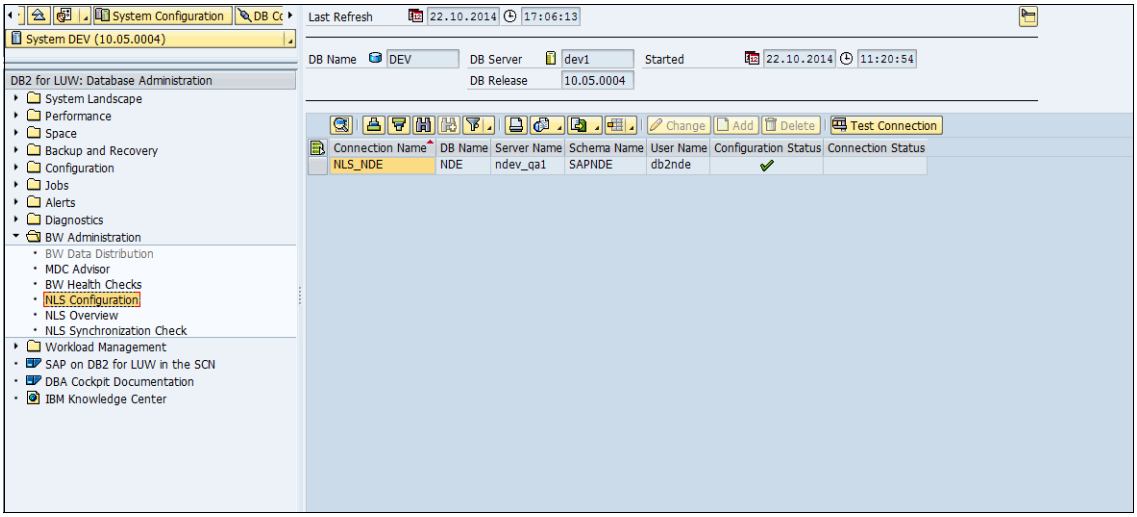

Figure 5-12 Administering the NLS database NDE in transaction DB02

From here, you are able to configure the NLS for use with InfoCubes and Data Store objects. To enable the DB2 BLU feature on these InfoProviders, set the Rsadmin Parameter DB6_NLS_USE_CDE to YES. Furthermore, you have to enable intrapartition as described in 4.2.4, “Enabling intrapartition parallelism for BLU” on page 51, and refer to SAP Note 1851832 for the DB2 parameter tuning on the NLS database.

For more information about the NLS database and the use of DB2 BLU, refer to section 8.9 of *Architecting and Developing DB2 with BLU Acceleration*, SG24-8212:

http://www.redbooks.ibm.com/abstracts/sg248212.html

Related publications

The publications listed in this section are considered particularly suitable for a more detailed discussion of the topics covered in this book.

IBM Redbooks

The following IBM Redbooks publications provide additional information about the topic in this document. Note that some publications referenced in this list might be available in softcopy only.

- *Architecting and Developing DB2 with BLU Acceleration*, SG24-8212
- *POWER processor, see Performance Optimization and Tuning Techniques for IBM Processors, including IBM POWER8*, SG24-8171-00
- *IBM PowerVM Virtualization Introduction and Configuration*, SG24-7940-05
- *Best Practices for DB2 on AIX 6.1 for Power Systems*, SG24-7821-00

You can search for, view, download or order these documents and other Redbooks, Redpapers, Web Docs, draft and additional materials, at the following website:

ibm.com/redbooks

Other publications

These publications are also relevant as further information sources:

- *Using AIX Active Memory Expansion with DB2*

  http://www-01.ibm.com/support/docview.wss?uid=swg21567898
- *SAP Note 1906519 and 1972803*
- SAP Database Administration with IBM DB2 by Faustmann, Greulich, Siegling, Wegener, Zimmermann Galileo Press/SAP Press

Online resources

These websites are also relevant as further information sources:

- DB2 Version 10.5 IBM Knowledge Center, Performance variables describes DB2 performance registry variables including DB2_RESOURCE_POLICY

  http://ibm.co/1M5xWWl
- The Star Schema Benchmark (SSB)

  http://www.cs.umb.edu/~poneil/StarSchemaB.PDF
- Server virtualization with IBM PowerVM

  http://www-03.ibm.com/systems/power/software/virtualization/resources.html

© Copyright IBM Corp. 2015. All rights reserved.

- POWER7 Virtualization Best Practice Guide
  http://ibm.co/1K5NLbj
- IBM AIX MPIO: Best practices and considerations
  http://www.ibm.com/developerworks/aix/library/au-aix-mpio
- AIX and VIOS Disk And Fibre Channel Adapter Queue Tuning
  http://www-03.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/TD105745
- IBM SAP Technical Brief: SAP on Power Systems Best Practices Guide
  http://ibm.co/1x1QcTi

Help from IBM

IBM Support and downloads

ibm.com/support

IBM Global Services

ibm.com/services

(0.1”spine)
0.1”<->0.169”
53<->89 pages

SG24-8258-00

ISBN 0738440671

Printed in U.S.A.

Get connected

Redbooks
ibm.com/redbooks