# SAP® Business Warehouse Performance and Space Comparison between the IBM® DB2® database and a competitive Database

Author: Andreas Christian
IBM SAP DB2 Center of Excellence

# Disclaimer & Trademarks

The information in this document may concern new products that IBM may or may not announce. Any discussion of OEM products is based upon information which has been publicly available and is subject to change. The specification of some of the features described in this presentation may change before the General Availability date of these products.
REFERENCES IN THIS PUBLICATION TO IBM PRODUCTS, PROGRAMS, OR SERVICES DO NOT IMPLY THAT IBM INTENDS TO MAKE THESE AVAILABLE IN ALL COUNTRIES IN WHICH IBM OPERATES.
IBM MAY HAVE PATENTS OR PENDING PATENT APPLICATIONS COVERING SUBJECT MATTER IN THIS DOCUMENT. THE FURNISHING OF THIS DOCUMENT DOES NOT IMPLY GIVING LICENSE TO THESE PATENTS.
TRADEMARKS
The following terms are registered trademarks of International Business Machines Corporation in the United States and/ or other countries: AIX, AIXwindows, AS/ 400, DB2, e( logo), IBM, IBM( logo), Information Warehouse, Netfinity, NUMA- Q, OS/ 2, OS/ 390, OS/ 400, Parallel Sysplex, PowerPC, PowerPC( logo), RISC System/ 6000, RS/ 6000, S/ 390, Sequent, SP2, System/ 390, The Engines of e-business, ThinkPad, Tivoli( logo), TURBOWAYS, VisualAge, WebSphere.
The following terms are trademarks of International Business Machines Corporation in the United States and/ or other countries: AIX/ L, AIX/ L( logo), AS/ 400e, DB2 OLAP Server, DB2 Universal Database, e- business (logo), HACMP/ 6000, Intelligent Miner, iSeries, Network Station, UMACenter, PowerPC Architecture, PowerPC 604, POWER2 Architecture, pSeries, Shark, SP, Tivoli Enterprise, TME 10, Videocharger, Visualization Data Explorer, xSeries, zSeries.
A full list of U. S. trademarks owned by IBM may be found at
http:// iplswww.nas.ibm.com/wpts/trademarks/trademar.htm
NetView, Tivoli and TME are registered trademarks and TME Enterprise is a trademark of Tivoli Systems, Inc. in the United States and/ or other countries.
Microsoft, Windows, Windows NT and the Windows logo are registered trademarks of Microsoft Corporation in the United States and/ or other countries.
SAP and related names like SAP NetWeaver are registered trademarks of SAP AG.
UNIX is a registered trademark in the United States and other countries licensed exclusively through The Open Group.
Oracle is a registered trademark of Oracle Corporation in the United Status and/or other countries.
LINUX is a registered trademark of Linus Torvalds.
Intel and Pentium are registered trademarks and MMX, Itanium, Pentium II Xeon and Pentium III Xeon are trademarks of Intel Corporation in the United States and/ or other countries.
Java and all Java- based trademarks and logos are trademarks of Sun Microsystems, Inc. in the United States and/ or other countries.
Other company, product and service names may be trademarks or service marks of others.

| **Document Version** (starting with 1.0) | **Status** (Draft/Review copy/ Released) | **Date** (DD.MM.YYYY) |
|---|---|---|
| 1.0 | Draft | 20.12.2007 |
| 1.1 | Review Copy | 07.01.2008 |
| 1.2 | Released | 08.01.2008 |
| | | |

# 1 Introduction

This document describes an IBM® DB2® Proof of Concept which was performed from July to October 2007 for a european customer with SAP® Business Warehouse. The project started with a DB2 compression assessment on HP Itanium® hardware which was the hardware plattform in use by the customer. It was then further extended by a performance benchmark on IBM System p hardware. The major goals of the assessment were as follows:

- Test DB2 deep compression with SAP BW and prove that **with DB2 9 deep compression at least 50% storage savings for database tables** can be achieved compared to the database system in use by the customer.
- Test performance of SAP BW on DB2 for dataload and queries. **Improve database performance by at least 20% with DB2 compared to the database system in use by the customer**.

# 2 Management Summary

Starting in July 2007, the *IBM SAP DB2 Center of Excellence* provided 6 weeks support for this DB2 Proof of Concept. The following tasks were performed by the *IBM SAP DB2 CoE Team*:

- Definition of the DB2 database configuration and selection of appropriate tables to be compressed.
- Migration of the SAP BW source system to the SAP BW/DB2 target systems (testsystem) on HP Hardware and IBM System p Hardware.
- Monitoring and performance tuning of the SAP BW/DB2 testsystems during the performance tests.

The project was finished in time and the defined KPI's were reached as follows:

- **Compression test on HP Itanium Hardware:** **60%** less space consumption for DB2 database tables and **42,7%** reduction of total database size compared to competitive database plattform on identical HP Itanium hardware.
- **Performance test on IBM System p Hardware:** In average factor **1.5** faster query response times on DB2 with filled database cache and factor **4.0** faster query response times with empty database cache compared to competitive database plattform on identical IBM System p hardware.

# 3 Detailed Test Results

## *3.1 Compression Test on HP Itanium Hardware*

| Used Space | Other Database | DB2 |
|---|---|---|
| Tables | 1.632 GB | 655 GB |
| Indexes | 673 GB | 683 GB |
| **Total** | **2.305 GB** | **1338 GB** |

DB2 required 42,7% less space than the competitive database.
DB2 tables required ~60% less space than the tables in the competitive database.

Roughly 600 of the largest database tables were compressed.

## 3.2 Performance Tests on IBM System p Hardware

The SAP BW query performance heavily depends on the database cache hitratio. The SAP BW query performance test comprised the following steps:

- The database was restarted to clean up the database cache.
- All test queries were executed sequentially in a defined order. This test case shows the response time with empty database cache. In this case all data was read from disk (each query accessed different tables).
- All test queries were sequentially executed a second time. This second execution shows the response time with filled database cache.

In a production environment multiple queries are executed concurrently. The database cache must be shared among all of these queries. Therefore, the actual cache hitratios in a production environment will be in between these two extreme cases ("empty cache" and "all selected data in cache"). The following table shows the query execution times with empty and with filled database cache.

| | **Runtime** (Empty Cache) | | **Runtime** (Filled Cache) | | |
|---|---|---|---|---|---|
| **SAP BW Report** | **Other DB** | **DB2** | **Other DB** | **DB2** | **#records selected / transfered** |
| **Q1: Management Reporting** | 76,18 | 5,79 | 19,73 | 5,75 | 748.568<br>72.288 |
| **Q2: CO-PA** | 6,15 | 6,92 | 1,03 | 0,97 | 23.256<br>4.489 |
| **Q3: Marketing & Sales** | 116,41 | 41,38 | 40,19 | 36,83 | 3.704.033<br>1.009.419 |
| **Q4: Inventory Management** | 14,65 | 6,03 | 6,16 | 5,20 | 932.001<br>797.330 |
| **Q5: Inventory Management** | 42,34 | 23,25 | 1,85 | 1,46 | 88.690<br>1.058 |
| **Q6: Purchasing Management** | 163,38 | 26,20 | 30,29 | 16,31 | 2.317.493<br>417.805 |
| **Q7: APO** | 30,29 | 2,64 | 2,11 | 0,92 | 31.690<br>2.033 |
| **Sum** | **449,41** | **112,22** | **101,36** | **67,43** | |

**Table 1: Comparison of Query Runtimes (seconds)**

If we compare the total execution times of DB2 and the competitive database we see significant performance advantages with DB2 in both test cases:

- Test with empty database cache: DB2 factor 4 faster
- Test with filled database cache: DB2 factor 1.5 faster

In the first test case with empty database cache there was a noticeable performance improvement with DB2. This test scenario is I/O bound because all qualifying pages must be read from disk. Since we used the DB2 compression feature for the large fact tables, the amount of data that had to be read from disc was drastically reduced.

The following figure provides on overview of the results of this test case.

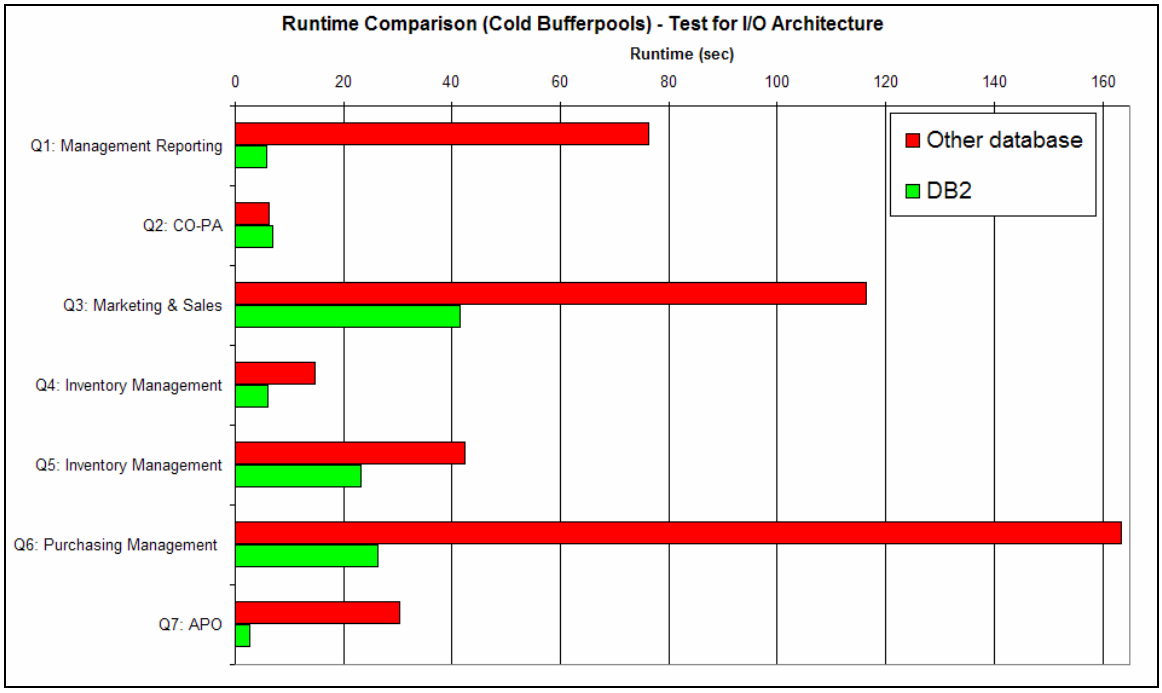


**Figure 1: Test with empty database cache**

In the second test case with filled database cache the performance improvement with DB2 is caused by DB2's inter-partition parallelism. By distributing the large fact tables across 8 partitions, a single SAP BW query is processed concurrently by 8 database agents.
The scaleout effect on DB2 becomes even more visible on larger servers with more than 8 CPUs/cores and more database partitions. DB2's shared nothing approach provides almost linear scalability for large complex SAP BW queries. This is a unique advantage of DB2's shared nothing architecture as opposed to the shared disk approach of other parallel databases.

Those queries which select many records on the database and do a lot of aggregation (i.e. return a small fraction of the selected data) gain most from

DB2 inter-partition parallelism. Good examples for this are the queries Q1 and Q6. The following figure provides on overview of the results of this test case.

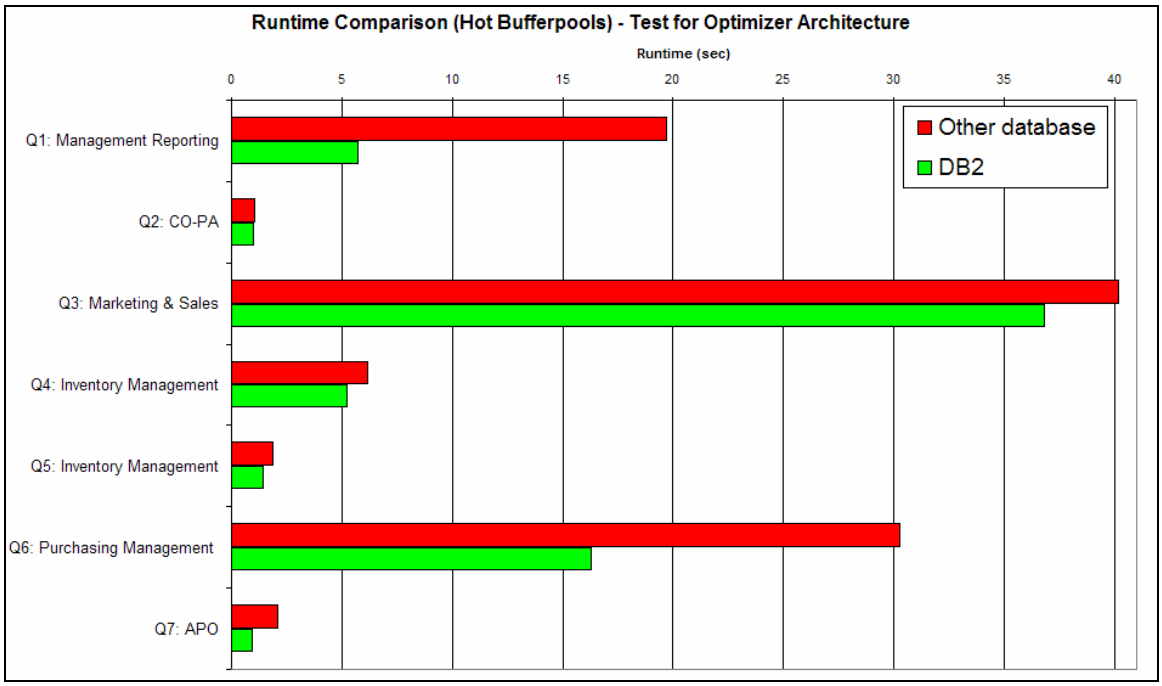


**Figure 2: Test with filled database cache**

Figure 3 shows the principle of DB2's inter-partition parallelism and shared nothing approach. Each SQL statement is split into sub sections which are processed in parallel on each of the partitions. The data is distributed across the partitions using hash partitioning. In SAP BW appropriate partition keys are automatically created on InfoCube fact tables, ODS tables, and PSA when these objects are activated. The SAP BW administrator controls the degree of parallelism by assigning a table to a SAP data class which refers to a distributed DB2 table space. The more partitions this table space spans, the more agents are working on the table whenever a query is executed.

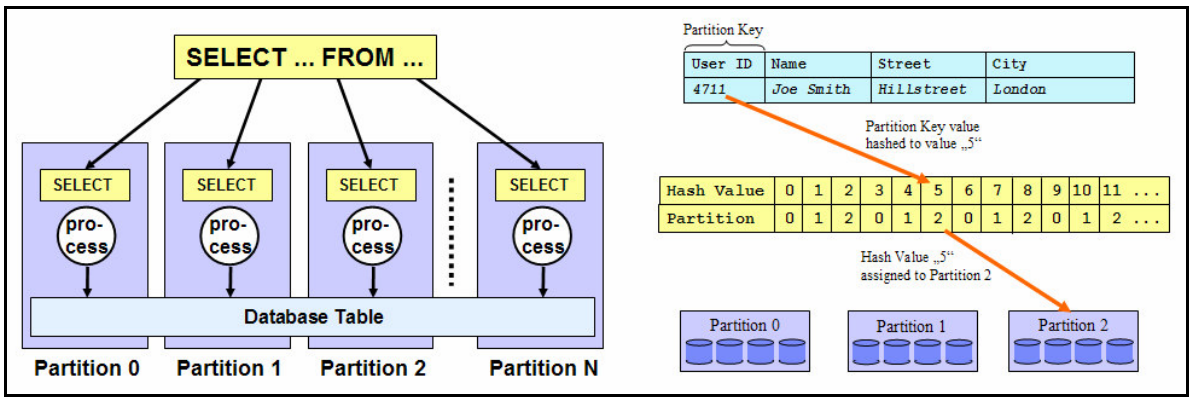


**Figure 3: DB2 inter-partition parallelism and hash partitioning**

# 4 IT Infrastructure

## 4.1 Database Servers

**Tests on HP Itanium Hardware:**
Server: HP RX6600 4 CPU's (2 cores each), 32GBRAM
Operating System: HPUX 11.23

**Tests on IBM System p Hardware:**
Server: p570 POWER6 (9117-MMA), 8x 4.7 GHz, 64 GB DDR2 667 MHz Memory
Operating System: AIX 5L V5.3 TL6 SP1

## 4.2 Storage

HP EVA 8000, 152 Disks, 4GB Cache, 3.5TB.
Server to storage subsystem connection via 2 Fibre Channel adapters.

## 4.3 Network

User-LAN 100BT, Service-LAN1000BT

## 4.4 Database Releases

IBM DB2 UDB 9 Fixpack 3

## 4.5 SAP BW Test System Configuration

| | |
|---|---|
| SAP Topology: | 2tier |
| | |
| SAP Version and conponents: | |
| SAP Basis Release | 7.00 SAPKB70011 |
| SAP BW Release | 7.00 |
| SAP Kernel-Version: | 7.00 Patchlevel 95 |

# 5 DB2 Configuration

## 5.1 DB2 Database Layout

Figure 4 shows the database layout which was used for the query performance tests. We used a total of 10 database partitions. The large fact- and ODS tables were distributed across 8 partitions to make sure that a single SAP BW query deploys all of the 8 cores that were configured in the server machine.
Partition 0 did not contain any large fact- or ODS tables to improve bufferpool quality on partition 0 and to allow for fast backup and restore. Since partition 0 contains smaller tables and some of the data is frequently accessed, this data should have a bufferpool assigned which is seperated from the large fact- and ODS data. Since partition 0 always needs to backed up first, this partition should be kept small to improve the backup performance.
A total of 5 GB memory was allocated for bufferpools. The bufferpool on partition 0 was roughly ten times larger than the bufferpools on partition 2 to 9. This configuration was used, because the number of logical reads on partition 0 was much larger than on the other partitions. Some of the test queries performed full table scans on master data tables which are located on partition 0.

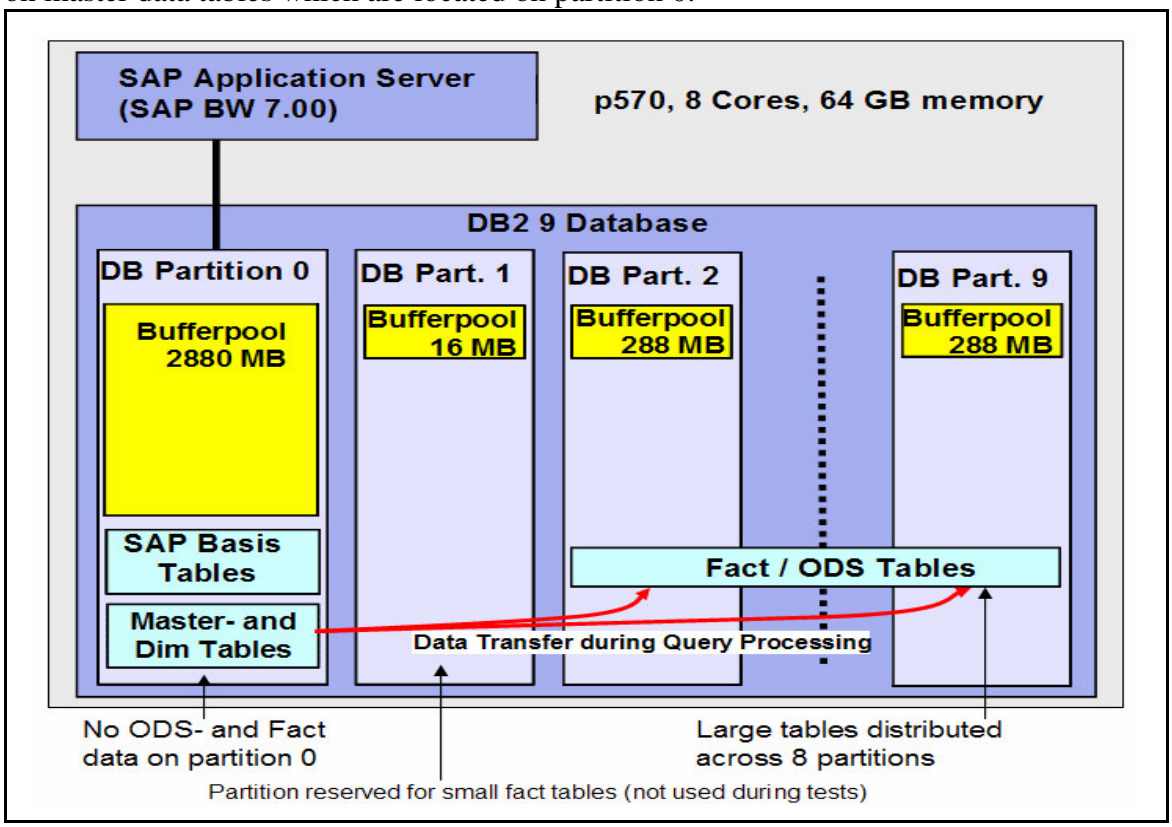


**Figure 4: DB2 Database Layout and SAP system configuration**

## 5.2 Storage and Filesystem Layout

Figure 5 shows the storage layout that was used for the tests on DB2. The available 152 disks were divided into two sets:

- We used 120 disks for data and indexes. 20 volume groups of equal size were created for data and indexes. On each of them one RAID5 volume and one sapdata filesystem was created.
- The remaining 32 disks were used for logging and were configured with RAID1.

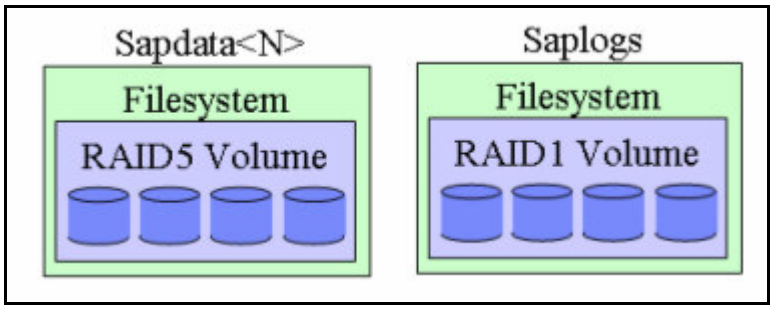


**Figure 5: Storage Layout**

## 5.3 Table Space Layout and Attributes

Fact and ODS table spaces were distributed across 8 database partitions to be able to deploy all 8 cores for a single SAP BW query. The following listing shows the layout of the distributed table spaces.

```
create regular  tablespace XCB#FACTD in nodegroup NGRP_2_TO_9_XCB
pagesize 16k managed by database
using (
   FILE '/db2/XCB/sapdata1/NODE0002/XCB#FACTD.container001' 1000 M ,
   FILE '/db2/XCB/sapdata2/NODE0002/XCB#FACTD.container002' 1000 M ,
   FILE '/db2/XCB/sapdata3/NODE0002/XCB#FACTD.container003' 1000 M ,
   FILE '/db2/XCB/sapdata4/NODE0002/XCB#FACTD.container004' 1000 M ,
   FILE '/db2/XCB/sapdata5/NODE0002/XCB#FACTD.container005' 1000 M
) on node ( 2 )  using (
   FILE '/db2/XCB/sapdata6/NODE0003/XCB#FACTD.container006' 1000 M ,
   FILE '/db2/XCB/sapdata7/NODE0003/XCB#FACTD.container007' 1000 M ,
   FILE '/db2/XCB/sapdata8/NODE0003/XCB#FACTD.container008' 1000 M ,
   FILE '/db2/XCB/sapdata9/NODE0003/XCB#FACTD.container009' 1000 M ,
   FILE '/db2/XCB/sapdata10/NODE0003/XCB#FACTD.container010' 1000 M
) on node ( 3 )  using (
   FILE '/db2/XCB/sapdata11/NODE0004/XCB#FACTD.container011' 1000 M ,
   FILE '/db2/XCB/sapdata12/NODE0004/XCB#FACTD.container012' 1000 M ,
   FILE '/db2/XCB/sapdata13/NODE0004/XCB#FACTD.container013' 1000 M ,
   FILE '/db2/XCB/sapdata14/NODE0004/XCB#FACTD.container014' 1000 M ,
   FILE '/db2/XCB/sapdata15/NODE0004/XCB#FACTD.container015' 1000 M
) on node ( 4 )  using (
   FILE '/db2/XCB/sapdata16/NODE0005/XCB#FACTD.container016' 1000 M ,
   FILE '/db2/XCB/sapdata17/NODE0005/XCB#FACTD.container017' 1000 M ,
   FILE '/db2/XCB/sapdata18/NODE0005/XCB#FACTD.container018' 1000 M ,
   FILE '/db2/XCB/sapdata19/NODE0005/XCB#FACTD.container019' 1000 M ,
   FILE '/db2/XCB/sapdata20/NODE0005/XCB#FACTD.container020' 1000 M
) on node ( 5 )  using (
   FILE '/db2/XCB/sapdata1/NODE0006/XCB#FACTD.container021' 1000 M ,
   FILE '/db2/XCB/sapdata2/NODE0006/XCB#FACTD.container022' 1000 M ,
   FILE '/db2/XCB/sapdata3/NODE0006/XCB#FACTD.container023' 1000 M ,
   FILE '/db2/XCB/sapdata4/NODE0006/XCB#FACTD.container024' 1000 M ,
   FILE '/db2/XCB/sapdata5/NODE0006/XCB#FACTD.container025' 1000 M
) on node ( 6 )  using (
   FILE '/db2/XCB/sapdata6/NODE0007/XCB#FACTD.container026' 1000 M ,
   FILE '/db2/XCB/sapdata7/NODE0007/XCB#FACTD.container027' 1000 M ,
   FILE '/db2/XCB/sapdata8/NODE0007/XCB#FACTD.container028' 1000 M ,
   FILE '/db2/XCB/sapdata9/NODE0007/XCB#FACTD.container029' 1000 M ,
   FILE '/db2/XCB/sapdata10/NODE0007/XCB#FACTD.container030' 1000 M
) on node ( 7 )  using (
   FILE '/db2/XCB/sapdata11/NODE0008/XCB#FACTD.container031' 1000 M ,
   FILE '/db2/XCB/sapdata12/NODE0008/XCB#FACTD.container032' 1000 M ,
   FILE '/db2/XCB/sapdata13/NODE0008/XCB#FACTD.container033' 1000 M ,
   FILE '/db2/XCB/sapdata14/NODE0008/XCB#FACTD.container034' 1000 M ,
   FILE '/db2/XCB/sapdata15/NODE0008/XCB#FACTD.container035' 1000 M
) on node ( 8 )  using (
   FILE '/db2/XCB/sapdata16/NODE0009/XCB#FACTD.container036' 1000 M ,
   FILE '/db2/XCB/sapdata17/NODE0009/XCB#FACTD.container037' 1000 M ,
   FILE '/db2/XCB/sapdata18/NODE0009/XCB#FACTD.container038' 1000 M ,
   FILE '/db2/XCB/sapdata19/NODE0009/XCB#FACTD.container039' 1000 M ,
   FILE '/db2/XCB/sapdata20/NODE0009/XCB#FACTD.container040' 1000 M
) on node ( 9 )
extentsize 2 prefetchsize automatic  dropped table recovery off  autoresize yes maxsize none
NO FILE SYSTEM CACHING
```

**Figure 6: Layout of distributed table spaces for fact and ODS data**

All other table spaces (standard R/3 tables, master tables, and dimension tables) were located on database partition 0. The following listing shows the layout of the table spaces that were not distributed.

```
create regular  tablespace XCB#STABD in nodegroup SAPNODEGRP_XCB
pagesize 16k managed by database
using (
   FILE '/db2/XCB/sapdata1/NODE0000/XCB#STABD.container000' 808 M ,
   FILE '/db2/XCB/sapdata2/NODE0000/XCB#STABD.container001' 808 M ,
   FILE '/db2/XCB/sapdata3/NODE0000/XCB#STABD.container002' 808 M ,
   FILE '/db2/XCB/sapdata4/NODE0000/XCB#STABD.container003' 808 M ,
   FILE '/db2/XCB/sapdata5/NODE0000/XCB#STABD.container004' 808 M ,
   FILE '/db2/XCB/sapdata6/NODE0000/XCB#STABD.container005' 808 M ,
   FILE '/db2/XCB/sapdata7/NODE0000/XCB#STABD.container006' 808 M ,
   FILE '/db2/XCB/sapdata8/NODE0000/XCB#STABD.container007' 808 M ,
   FILE '/db2/XCB/sapdata9/NODE0000/XCB#STABD.container008' 808 M ,
   FILE '/db2/XCB/sapdata10/NODE0000/XCB#STABD.container009' 808 M ,
   FILE '/db2/XCB/sapdata11/NODE0000/XCB#STABD.container010' 808 M ,
   FILE '/db2/XCB/sapdata12/NODE0000/XCB#STABD.container011' 808 M ,
   FILE '/db2/XCB/sapdata13/NODE0000/XCB#STABD.container012' 808 M ,
   FILE '/db2/XCB/sapdata14/NODE0000/XCB#STABD.container013' 808 M ,
   FILE '/db2/XCB/sapdata15/NODE0000/XCB#STABD.container014' 808 M ,
   FILE '/db2/XCB/sapdata16/NODE0000/XCB#STABD.container015' 808 M ,
   FILE '/db2/XCB/sapdata17/NODE0000/XCB#STABD.container016' 808 M ,
   FILE '/db2/XCB/sapdata18/NODE0000/XCB#STABD.container017' 808 M ,
   FILE '/db2/XCB/sapdata19/NODE0000/XCB#STABD.container018' 808 M ,
   FILE '/db2/XCB/sapdata20/NODE0000/XCB#STABD.container019' 808 M )
  on node ( 0 )
extentsize 2 prefetchsize automatic  dropped table recovery off  autoresize yes maxsize none
NO FILE SYSTEM CACHING
```

**Figure 7: Layout of non-distributed table spaces**

The following table space attributes were used:

- Type: Database managed using file system
- Page size: 16K
- Extent size: 2
- Prefetch size: AUTOMATIC
- Dropped table recovery: OFF
- Autoresize: YES
- Maxsize: NONE
- File system caching: NO

## 5.4 Database Parameters

Most of the configuration parameters were set according to SAP note 899322. The most important parameters are listed below together with a short explanation.

DB2 Database Configuration Parameters

```
SHEAPTHRES_SHR = 100000 (default 40000)
      This value resulted in roughly 1 GB memory for sortheap for the
      database (5 GB were allocated for bufferpools).
SORTHEAP = 25000 (default 15000)
      Sortheap was larger than default to reduce amount of sort
      overflows.
LOCKLIST = 40000 (default)
CATALOGCACHE_SZ = 10000 (default 2560)
      Catalog cache was larger than default to increase catalog cache
      hitratio.
PCKCACHESZ = 5120 (default)
DBHEAP = 35000 (default)
AUTO_RUNSTATS = OFF
      Auto runstats was switched off to avoid collection of
      statistics during measurements.
DFT_QUERYOPT = 5
      For the SAP Queries we used opt level 7.
```

DB2 Registry Settings

```
DB2_REDUCED_OPTIMIZATION=4,INDEX,JOIN,NO_HSJN_BUILD_FACT,STARJN_CARD_
SKEW,NO_SORT_MGJOIN,CART OFF [O]
      The NO_TQ_FACT setting was removed from the values to reduce
      communication overhead between the nodes.
DB2_PARALLEL_IO=*
    For each table space not more than one container was placed on a
    RAID5 volume. Therefore we enabled parallel I/O to have multiple
    I/O servers working on each container.
```

# 6 Appendix

## 6.1 DB2 Database Configuration Parameters

```
      Database Configuration for Database xcb

 Database configuration release level                    = 0x0b00
 Database release level                                  = 0x0b00

 Database territory                                      = en_US
 Database code page                                      = 1208
 Database code set                                       = UTF-8
 Database country/region code                            = 1
 Database collating sequence                             =
IDENTITY_16BIT
 Alternate collating sequence              (ALT_COLLATE) =
 Database page size                                      = 16384

 Dynamic SQL Query management           (DYN_QUERY_MGMT) = DISABLE

 Discovery support for this database       (DISCOVER_DB) = ENABLE

 Restrict access                                         = NO
 Default query optimization class         (DFT_QUERYOPT) = 5
 Degree of parallelism                      (DFT_DEGREE) = 1
 Continue upon arithmetic exceptions   (DFT_SQLMATHWARN) = NO
 Default refresh age                  (DFT_REFRESH_AGE) = 0
 Default maintained table types for opt (DFT_MTTB_TYPES) = SYSTEM
 Number of frequent values retained     (NUM_FREQVALUES) = 10
 Number of quantiles retained            (NUM_QUANTILES) = 20

 Backup pending                                          = NO

 Database is consistent                                  = NO
 Rollforward pending                                     = NO
 Restore pending                                         = NO

 Multi-page file allocation enabled                      = YES

 Log retain for recovery status                          = NO
 User exit for logging status                            = YES

 Self tuning memory                    (SELF_TUNING_MEM) = OFF
 Size of database shared memory (4KB)  (DATABASE_MEMORY) = COMPUTED
 Database memory threshold               (DB_MEM_THRESH) = 10
 Max storage for lock list (4KB)               (LOCKLIST) = 40000
 Percent. of lock lists per application        (MAXLOCKS) = 90
 Package cache size (4KB)                   (PCKCACHESZ) = 5120
 Sort heap thres for shared sorts (4KB) (SHEAPTHRES_SHR) = 100000
 Sort list heap (4KB)                          (SORTHEAP) = 25000
```

```
 Database heap (4KB)                                (DBHEAP) = 35000
 Catalog cache size (4KB)                (CATALOGCACHE_SZ) = 10000
 Log buffer size (4KB)                          (LOGBUFSZ) = 1024
 Utilities heap size (4KB)                   (UTIL_HEAP_SZ) = 5000
 Buffer pool size (pages)                       (BUFFPAGE) = 10000
 Max size of appl. group mem set (4KB) (APPGROUP_MEM_SZ) = 128000
 Percent of mem for appl. group heap   (GROUPHEAP_RATIO) = 25
 Max appl. control heap size (4KB)     (APP_CTL_HEAP_SZ) = 1600

 SQL statement heap (4KB)                        (STMTHEAP) = 5120
 Default application heap (4KB)                (APPLHEAPSZ) = 3072
 Statistics heap size (4KB)                 (STAT_HEAP_SZ) = 15000

 Interval for checking deadlock (ms)           (DLCHKTIME) = 10000
 Lock timeout (sec)                           (LOCKTIMEOUT) = 3600

 Changed pages threshold                  (CHNGPGS_THRESH) = 40
 Number of asynchronous page cleaners     (NUM_IOCLEANERS) = AUTOMATIC
 Number of I/O servers                     (NUM_IOSERVERS) = AUTOMATIC
 Index sort flag                              (INDEXSORT) = YES
 Sequential detect flag                       (SEQDETECT) = YES
 Default prefetch size (pages)          (DFT_PREFETCH_SZ) = 32

 Track modified pages                           (TRACKMOD) = ON

 Default number of containers                              = 1
 Default tablespace extentsize (pages)    (DFT_EXTENT_SZ) = 2

 Max number of active applications              (MAXAPPLS) = AUTOMATIC
 Average number of active applications         (AVG_APPLS) = AUTOMATIC
 Max DB files open per application              (MAXFILOP) = 1950

 Log file size (4KB)                            (LOGFILSIZ) = 16380
 Number of primary log files                   (LOGPRIMARY) = 20
 Number of secondary log files                  (LOGSECOND) = 40
 Changed path to log files                     (NEWLOGPATH) =
 Path to log files                                          =
/db2/XCB/log_dir/NODE0000/
 Overflow log path                       (OVERFLOWLOGPATH) =
 Mirror log path                           (MIRRORLOGPATH) =
 First active log file                                      =
S0000008.LOG
 Block log on disk full                  (BLK_LOG_DSK_FUL) = YES
 Percent max primary log space by transaction  (MAX_LOG) = 0
 Num. of active log files for 1 active UOW(NUM_LOG_SPAN) = 0

 Group commit count                             (MINCOMMIT) = 1
 Percent log file reclaimed before soft chckpt  (SOFTMAX) = 300
 Log retain for recovery enabled                (LOGRETAIN) = OFF
 User exit for logging enabled                   (USEREXIT) = OFF

 HADR database role                                         = STANDARD
 HADR local host name                    (HADR_LOCAL_HOST) =
 HADR local service name                  (HADR_LOCAL_SVC) =
 HADR remote host name                 (HADR_REMOTE_HOST) =
 HADR remote service name               (HADR_REMOTE_SVC) =
 HADR instance name of remote server   (HADR_REMOTE_INST) =
 HADR timeout value                        (HADR_TIMEOUT) = 120
 HADR log write synchronization mode      (HADR_SYNCMODE) = NEARSYNC
```

```
 First log archive method                 (LOGARCHMETH1) =
DISK:/db2/XCB/LogArchive/
 Options for logarchmeth1                  (LOGARCHOPT1) =
 Second log archive method                (LOGARCHMETH2) = OFF
 Options for logarchmeth2                  (LOGARCHOPT2) =
 Failover log archive path                (FAILARCHPATH) =
 Number of log archive retries on error   (NUMARCHRETRY) = 5
 Log archive retry Delay (secs)         (ARCHRETRYDELAY) = 20
 Vendor options                               (VENDOROPT) =

 Auto restart enabled                      (AUTORESTART) = ON
 Index re-creation time and redo index build  (INDEXREC) = SYSTEM
(RESTART)
 Log pages during index build            (LOGINDEXBUILD) = OFF
 Default number of loadrec sessions    (DFT_LOADREC_SES) = 1
 Number of database backups to retain   (NUM_DB_BACKUPS) = 12
 Recovery history retention (days)     (REC_HIS_RETENTN) = 60

 TSM management class                    (TSM_MGMTCLASS) =
 TSM node name                            (TSM_NODENAME) =
 TSM owner                                   (TSM_OWNER) =
 TSM password                             (TSM_PASSWORD) =

 Automatic maintenance                       (AUTO_MAINT) = OFF
   Automatic database backup              (AUTO_DB_BACKUP) = OFF
   Automatic table maintenance            (AUTO_TBL_MAINT) = OFF
     Automatic runstats                    (AUTO_RUNSTATS) = OFF
     Automatic statistics profiling     (AUTO_STATS_PROF) = OFF
       Automatic profile updates          (AUTO_PROF_UPD) = OFF
     Automatic reorganization                (AUTO_REORG) = OFF
```

## 6.2 DB2 Database Manager Parameters

```
Database Manager Configuration

 Node type = Enterprise Server Edition with local and remote clients
 Database manager configuration release level            = 0x0b00
 CPU speed (millisec/instruction)             (CPUSPEED) = 2.834e-07
 Communications bandwidth (MB/sec)      (COMM_BANDWIDTH) = 1.000e+02

 Max number of concurrently active databases     (NUMDB) = 8
 Federated Database System Support           (FEDERATED) = NO
 Transaction processor monitor name        (TP_MON_NAME) =

 Default charge-back account           (DFT_ACCOUNT_STR) =

 Java Development Kit installation path       (JDK_PATH) =
/db2/db2xcb/sqllib/java/jdk64

 Diagnostic error capture level              (DIAGLEVEL) = 3
 Notify Level                               (NOTIFYLEVEL) = 3
 Diagnostic data directory path                (DIAGPATH) =
/db2/XCB/db2dump

 Default database monitor switches
   Buffer pool                          (DFT_MON_BUFPOOL) = ON
   Lock                                    (DFT_MON_LOCK) = ON
   Sort                                    (DFT_MON_SORT) = ON
   Statement                               (DFT_MON_STMT) = ON
   Table                                  (DFT_MON_TABLE) = ON
   Timestamp                         (DFT_MON_TIMESTAMP) = ON
   Unit of work                             (DFT_MON_UOW) = ON
 Monitor health of instance and databases   (HEALTH_MON) = OFF

 SYSADM group name                         (SYSADM_GROUP) = DBXCBADM
 SYSCTRL group name                       (SYSCTRL_GROUP) = DBXCBCTL
 SYSMAINT group name                     (SYSMAINT_GROUP) = DBXCBMNT
 SYSMON group name                         (SYSMON_GROUP) =

 Client Userid-Password Plugin          (CLNT_PW_PLUGIN) =
 Client Kerberos Plugin                (CLNT_KRB_PLUGIN) =
 Group Plugin                             (GROUP_PLUGIN) =
 GSS Plugin for Local Authorization    (LOCAL_GSSPLUGIN) =
 Server Plugin Mode                    (SRV_PLUGIN_MODE) = UNFENCED
 Server List of GSS Plugins      (SRVCON_GSSPLUGIN_LIST) =
 Server Userid-Password Plugin        (SRVCON_PW_PLUGIN) =
 Server Connection Authentication          (SRVCON_AUTH) =
NOT_SPECIFIED
 Database manager authentication        (AUTHENTICATION) =
SERVER_ENCRYPT
 Cataloging allowed without authority   (CATALOG_NOAUTH) = NO
 Trust all clients                      (TRUST_ALLCLNTS) = YES
 Trusted client authentication          (TRUST_CLNTAUTH) = CLIENT
 Bypass federated authentication            (FED_NOAUTH) = NO

 Default database path                       (DFTDBPATH) = /db2/XCB
```

```
 Database monitor heap size (4KB)           (MON_HEAP_SZ) = 128
 Java Virtual Machine heap size (4KB)      (JAVA_HEAP_SZ) = 2048
 Audit buffer size (4KB)                   (AUDIT_BUF_SZ) = 0
 Size of instance shared memory (4KB)  (INSTANCE_MEMORY) = AUTOMATIC
 Backup buffer default size (4KB)             (BACKBUFSZ) = 1024
 Restore buffer default size (4KB)            (RESTBUFSZ) = 1024

 Sort heap threshold (4KB)                   (SHEAPTHRES) = 100000

 Directory cache support                      (DIR_CACHE) = NO

 Application support layer heap size (4KB)   (ASLHEAPSZ) = 16
 Max requester I/O block size (bytes)         (RQRIOBLK) = 65000
 Query heap size (4KB)                    (QUERY_HEAP_SZ) = 2000

 Workload impact by throttled utilities(UTIL_IMPACT_LIM) = 10

 Priority of agents                           (AGENTPRI) = SYSTEM
 Max number of existing agents               (MAXAGENTS) = 1024
 Agent pool size                         (NUM_POOLAGENTS) = 10
 Initial number of agents in pool        (NUM_INITAGENTS) = 5
 Max number of coordinating agents      (MAX_COORDAGENTS) = (MAXAGENTS
- NUM_INITAGENTS)
 Max no. of concurrent coordinating agents  (MAXCAGENTS) =
MAX_COORDAGENTS
 Max number of client connections       (MAX_CONNECTIONS) =
MAX_COORDAGENTS

 Keep fenced process                         (KEEPFENCED) = NO
 Number of pooled fenced processes          (FENCED_POOL) = 5
 Initial number of fenced processes      (NUM_INITFENCED) = 0

 Index re-creation time and redo index build  (INDEXREC) = RESTART

 Transaction manager database name          (TM_DATABASE) = 1ST_CONN
 Transaction resync interval (sec)      (RESYNC_INTERVAL) = 180

 SPM name                                     (SPM_NAME) =
 SPM log size                           (SPM_LOG_FILE_SZ) = 256
 SPM resync agent limit                  (SPM_MAX_RESYNC) = 20
 SPM log path                              (SPM_LOG_PATH) =

 TCP/IP Service name                          (SVCENAME) = sapdb2XCB
 Discovery mode                               (DISCOVER) = SEARCH
 Discover server instance                (DISCOVER_INST) = ENABLE

 Maximum query degree of parallelism    (MAX_QUERYDEGREE) = 1
 Enable intra-partition parallelism      (INTRA_PARALLEL) = NO

 Maximum Asynchronous TQs per query     (FEDERATED_ASYNC) = 0

 No. of int. communication buffers(4KB)(FCM_NUM_BUFFERS) = AUTOMATIC
 No. of int. communication channels   (FCM_NUM_CHANNELS) = AUTOMATIC
 Node connection elapse time (sec)          (CONN_ELAPSE) = 10
 Max number of node connection retries (MAX_CONNRETRIES) = 5
 Max time difference between nodes (min) (MAX_TIME_DIFF) = 60
 db2start/db2stop timeout (min)         (START_STOP_TIME) = 10
```

## 6.3 DB2 Registry Settings

```
DB2_OPT_MAX_TEMP_SIZE=10240 [DB2_WORKLOAD]
DB2_WORKLOAD=SAP
DB2_TRUNCATE_REUSESTORAGE=IMPORT [DB2_WORKLOAD]
DB2_MDC_ROLLOUT=YES [DB2_WORKLOAD]
DB2RSHCMD=/usr/bin/ssh
DB2_SKIPINSERTED=YES [DB2_WORKLOAD]
DB2_VIEW_REOPT_VALUES=YES [DB2_WORKLOAD]
DB2_OBJECT_TABLE_ENTRIES=65532 [DB2_WORKLOAD]
DB2_OPTPROFILE=YES [DB2_WORKLOAD]
DB2_IMPLICIT_UNICODE=YES [DB2_WORKLOAD]
DB2_INLIST_TO_NLJN=YES [DB2_WORKLOAD]
DB2_MINIMIZE_LISTPREFETCH=YES [DB2_WORKLOAD]
DB2_UPDATE_PART_KEY=YES [DB2_WORKLOAD]
DB2_REDUCED_OPTIMIZATION=4,INDEX,JOIN,NO_HSJN_BUILD_FACT,STARJN_CARD_
SKEW,NO_SORT_MGJOIN,CART OFF [O]
DB2NOTIFYVERBOSE=YES [DB2_WORKLOAD]
DB2_INTERESTING_KEYS=YES [DB2_WORKLOAD]
DB2_EVALUNCOMMITTED=YES [DB2_WORKLOAD]
DB2_ANTIJOIN=EXTEND [DB2_WORKLOAD]
DB2MEMMAXFREE=2000000 [DB2_WORKLOAD]
DB2ENVLIST=INSTHOME SAPSYSTEMNAME dbs_db6_schema DIR_LIBRARY
SHLIB_PATH
DB2_RR_TO_RS=YES [DB2_WORKLOAD]
DB2_DROP_NO_WAIT=YES [DB2_WORKLOAD]
DB2_FORCE_FCM_BP=YES [DB2_WORKLOAD]
DB2COUNTRY=1
DB2COMM=TCPIP [DB2_WORKLOAD]
DB2_PARALLEL_IO=*
```