
DB2 9 Row Compression in a SAP R/3 production system

Compression Attributes and Performance Influence

White Paper

Analysis Documentation

Version 1.2

IBM SAP DB2

Center of Excellence

Revision date: 01.03.2007 15:44

Author: Waldemar Gaida

Management Summary

This document describes the results of an analysis that was executed at customer site. The system investigated was a SAP R/3 production system using SAP R/3 Enterprise with the industry solution insurance (IS-Insurance).

The analysis comprises the effects on the compression of large tables using the DB2 9 row compression feature. Furthermore the effect on response times (Dialog and Batch) was investigated.

The database software was upgraded from DB2 V8 FP 12 to DB2 9 using the standard upgrade procedure. The tables to compress were determined by using the process described in SAP note 980067.

The results that we received are remarkable and very encouraging related to the usage of DB2 compression in a SAP R/3 system.

The compression of the chosen 310 tables resulted in an **average compression rate of 70%** (from 132,5 GB to 39,7 GB).

The **overall space saving of the database** was **43%** (from 266,1 GB to 151,1 GB), which also includes the effects of the reorganization and table movements to reduce table space high water marks. This was in spite of having the large RID attribute switched on and being activated during REORG for the selected 310 tables, which increased the size of the indexes.

Another main focus was the answer to the question how the compression would effect the **response time behavior** of the system. Taking as data basis one week before and after the compression the result was a **significant improvement** on performance of dialog and batch processing (both improved about **25%**).

If the systems space is fragmented, then the resulting effects will be caused by either the compression and the space saving during the reorganization of tables and indexes.

This investigation shows that a SAP R/3 system (OLTP), that was compressed using the standard recommendation of SAP according to SAP-Note 980067, can not only gain disk space but also can improve performance significantly.

Table of Contents

- 0 CHANGE MANAGEMENT 4**
- 1 DISCLAIMER & TRADEMARKS..... 5**
- 2 INTRODUCTION..... 6**
- 3 SAP SYSTEM ENVIRONMENT 7**
 - 3.1 SYSTEM ENVIRONMENT 7
- 4 DB2 ROW COMPRESSION – THEORY AND PRACTICAL USE 8**
 - 4.1 HOW TO ACTIVATE COMPRESSION..... 8
- 5 COMPRESS PROCEDURE 10**
 - 5.1 DB2 SOFTWARE UPGRADE 10
 - 5.2 IDENTIFYING COMPRESSION CANDIDATES 10
 - 5.3 TABLE SPACE HIGH-WATER-MARK (HWM) CONSIDERATIONS 10
- 6 RESULTS 12**
 - 6.1 COMPRESSION RATIOS 12
 - 6.2 LARGEST TABLES 18
 - 6.3 CPU USAGE 19
 - 6.4 RESPONSE TIMES 20
 - 6.5 DIALOG RESPONSE TIMES 20
 - 6.6 BATCH RESPONSE TIMES 23
- 7 APPENDIX..... 24**
 - 7.1 DEFINITIONS 24
 - 7.2 DB2 REGISTRY VARIABLES 24
 - 7.3 DBM CONFIGURATION 25
 - 7.4 DB CONFIGURATION..... 26
- 8 LIST OF LITERATURE..... 29**
- 9 TABLE OF FIGURES / TABLES / EQUATIONS 30**
 - 9.1 FIGURES..... 30
 - 9.2 TABLES 30
 - 9.3 EQUATIONS..... 30

0 Change Management

All important changes and additions for every version of this document are documented in this chapter.

Description of Changes

Version	Date	Description
1.1	16.02.2007	First version
1.2	01.03.2007	Addition of Disclaimer & Trademarks

1 Disclaimer & Trademarks

The information in this presentation 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, NUMACenter, 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 States 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.

2 Introduction

The DB2 9 row compression feature (sometimes also named deep compression) is capable of saving space in database tables. After having first results from tests in the labs, it was interesting to investigate this feature on a real life system of a customer. The main focus of this investigation besides the compression results concerning disk space was the influence of the compression on dialog and batch response times of a productive SAP R/3 system. This documentation describes the basic conditions of the system environment, the procedure of the switch to a DB2 9 compressed system and the measured results.

3 SAP System Environment

The SAP system was a SAP R/3 Enterprise (4.70 with Extension Set 1.10) and the industry solution IS-Insurance. The following figure shows the component versions:

COMPONENT	RELEASE	EXTRELEASE	COMP_TYPE
SAP_APPL	470	0025	R
SAP_BASIS	620	0060	S
SAP_ABA	620	0060	S
SAP_HR	470	0059	R
PI	2004_1_470	0012	P
INSURANCE	471	0016	I
ST-A/PI	01H_R3_470	0000	C
ST-PI	2005_1_620	0004	P
WPP-PI	500_46C	0003	C
EA-APPL	110	0021	N
EA-FINSERV	110	0021	N
EA-GLTRADE	110	0000000006	N
EA-HR	110	0030	N
EA-IPPE	110	0020	P
EA-PS	110	0000000006	N
EA-RETAIL	110	0000000006	N
PI_BASIS	2005_1_620	0008	P

Figure 1 – Components of the SAP system

3.1 System Environment

Server Model	IBM,7038-6M2
CPU	4 x PowerPC_RS64-III (1452 MHz)
RAM	17 GB
Interfaces	2 x FC
Disk Storage (SAN)	EMC Symmetrix
Operating System	AIX 5 64-bit
Database	DB2 9

Figure 2 - Productive System Environment

4 DB2 Row Compression – Theory and Practical Use

DB2 9 uses a row compression that is based on the LZ (Lempel-Ziv) algorithm [1].

It uses a dictionary to make references to often used patterns in a table.

The following figure shows how this works (taken from [2]):

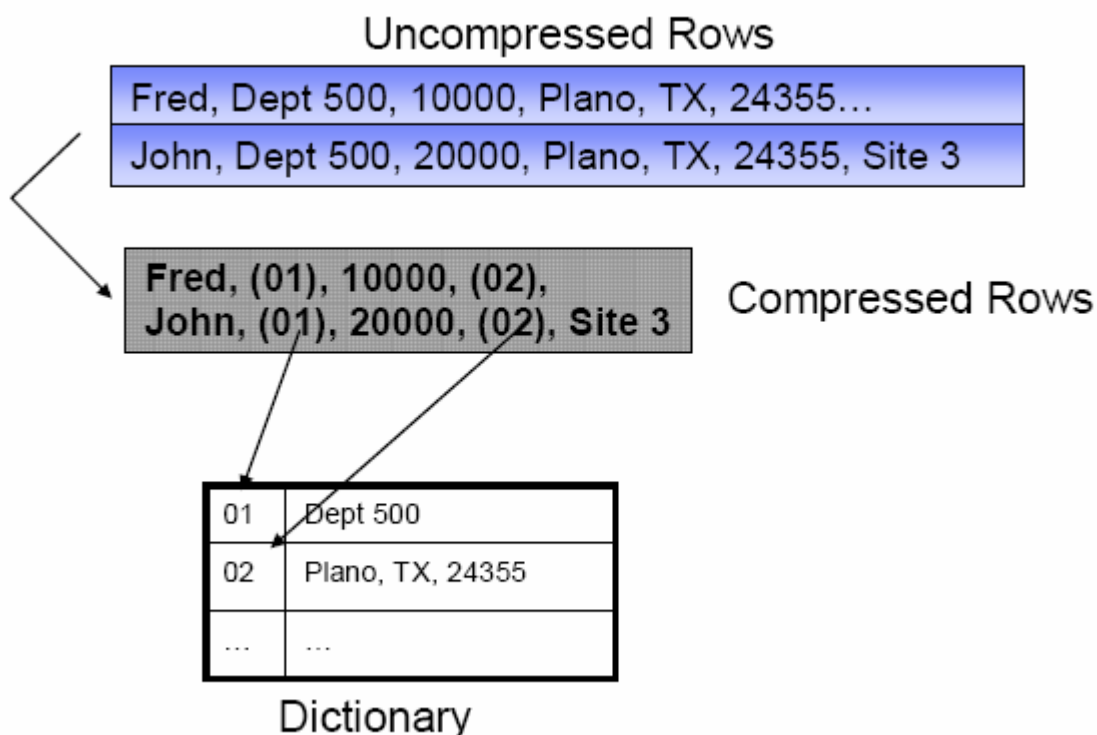


Figure 3 - Compression algorithm

Instead of the whole pattern only the 12bit key is stored in the compressed row. The dictionary contains patterns that are built up either of one column or of several adjacent columns.

4.1 How to Activate Compression

To activate the compression for a table <TAB1> you have to do the following steps:

```
ALTER TABLE <TAB1> COMPRESSION YES
```

This enables the table to use the compression feature. If the table is new you can use the compression attribute on the CREATE TABLE statement.

```
REORG TABLE <TAB1>
```

During the (offline) reorganization of the table the dictionary will be built up or updated and the data will be compressed according to the dictionary. Future insert, update, import or load activities will also take notice of the dictionary and compress the rows according to it.

To refresh the statistics for the compressed table it is useful to run the **RUNSTATS** command.

To estimate the compression rate of a table, the INSPECT command can be used in the following way:

```
INSPECT ROWCOMPESTIMATE TABLE NAME <table_name>  
RESULTS KEEP <file_name>
```

This will estimate the savings reachable by the compression. The result-file <file_name> is placed in the DIAGPATH and must be converted by the program `db2inspf` to a human-readable format.

If a table is enabled for compression, the INSPECT command will build up the compression dictionary while the table is “online”. To do this consistently the database will set an “Exclusive Table Alter Lock” and an “Intent on Exclusive Table Lock”. After the dictionary is built up all subsequent DML operations (like update, insert, load) will compress rows according to the compression dictionary.

This method can be valuable if a customer cannot afford a downtime that is long enough to do an offline reorg on the tables that should be compressed (e.g. in a 24 x 7 production environment with very limited maintenance time-windows).

5 Compress Procedure

5.1 DB2 Software Upgrade

To be able to use DB2 9 row compression, the database software had to be upgraded from DB2 V8 FP12 to DB2 9. The procedure followed the descriptions given in [3]. According to chapter 5.7 of this guide large RID support was enabled.

Due to the enabling of large RID support, the size of the indexes belonging to the tables that were reorganized increased. If the index is heavily fragmented before compression, the space gain due to reorganization may outweigh the increase due to large RID.

Find some interesting remarks and explanations on large RID support in [4]. This Blog also gives some hints on the estimation of the increase of indexes by activating large RID.

5.2 Identifying Compression Candidates

To find appropriate candidates for compression the **SAP-note 980067** (DB6: Using DB2 9 Row Compression) was used.

This SAP-Note delivers scripts and an ABAP-program. Either of these can be used to find candidates for compression. Since the ABAP-program was not available at the time the compression was done, the script-related procedure was used.

Important Hint

When using the script-based procedure be aware that you must not close your connection in between, because the results are put into SESSION-tables.

The procedure identified 310 candidates for compression. The details about the names and the compression ratios are described in chapter 6.1.

5.3 Table Space High-Water-Mark (HWM) Considerations

If you switch on DB2 9 Compression one of the goals normally is to save disk space. You can use this gained disk space in two different ways. You can use it to prolongue your time until you need to buy new disks because of database growth. Or you can free up some disks to use them with other applications.

If you want to go the second way, you must be able to give back disk space to the operating system. This is normally done by shrinking down a table space and deleting one or more table space containers.

A table space can only be shrunk down to its HWM. If there is freespace in the table space the HWM can be lowered by using db2dart and doing some table reorgs (see [5] pages 116 ff. and SAP-notes 152531, 486559). In some cases you can not lower the HWM of a table space. This is the case when the HWM is held by an object table.

If you are not able to lower the HWM, it might be feasible to move tables to another table space with DB6CONV. This could be done in the following steps:

1. Create new table space <TBS_N>
2. move table <TAB1> to new table space <TBS_N> with DB6CONV
3. enable table <TAB1> for compression (see 4.1)
4. reorganize table <TAB1> to build up dictionary and compress table
5. Repeat steps 2-4 for every table <TAB2> ... <TABn>

This sequential procedure ensures that there will be no fragmentation after the compression of the tables.

For a description on how to use DB6CONV look at [5] Chapter 4.6 (p. 124 ff.) and the SAP-notes 362325 (Using DB6CONV) and 817709 (DB6CONV-version history).

6 Results

In this chapter we will either have a look at the results in form of achieved compression ratios and we will compare performance data of the system before and after having carried out the compression procedure described in ch. 5.

6.1 Compression Ratios

The following table shows the compression ratios¹ of the tables that were selected for compression. The table is sorted descending by the compression ratio.

The values for the columns PCTPAGESSAVED and PCTROWSCOMPRESSED were taken from the view SYSCAT.TABLES.

PCTPAGESSAVED describes the amount of pages that were saved measured in percent.

PCTROWSCOMPRESSED describes the amount of rows that were compressed in percent.

TABNAME	PCTPAGESSAVED [%]	PCTROWSCOMPRESSED [%]
SE16N_CD_DATA	87	100,00
ZZKV2T	86	100,00
TPRI_PAR	86	100,00
ZZKV2P	85	100,00
ZZKV1T	85	100,00
ZZKV1P	85	100,00
ISSRFLDET	85	100,00
COSS	85	100,00
COEJ	85	100,00
ANLC	85	100,00
ZZISDET	84	100,00
PA0012	84	100,00
PA0008	84	100,00
COSP	84	100,00
COEJR	84	100,00
ANLP	84	100,00
ANLB	84	100,00
ZVTBFHAPO	83	100,00
ZUPDATEVZZKOPO	83	100,00
ZSAVEVZZKOPO_DEM	83	100,00
ZSAVEVZZKOPO	83	100,00
ZCORRVZZKOPO	83	100,00
SWWWIHEAD	83	100,00
EWUCOSUM	83	100,00
DFKKOPK	83	100,00
DB6PMSQ_DB	83	100,00
COSR	83	100,00

¹ For the definition see App. 7.1

VZZKOPO	82	100,00
VZZKOKO	82	100,00
VTBFINKO	82	100,00
T599U	82	99,95
LFC1	82	100,00
KNC1	82	100,00
BSID	82	100,00
ANEA	82	100,00
ZZKV2A	81	100,00
ZZKV1A	81	100,00
VVSCPOS	81	100,00
T512W	81	100,00
SWD_MNODES	81	100,00
PA0079	81	100,00
PA0013	81	100,00
DFKKOP	81	100,00
ZZKV1O	80	100,00
ZZISDEO	80	100,00
VVSCITEM	80	100,00
TRQT_FLOWGROUP	80	100,00
PA9637	80	100,00
PA0014	80	100,00
PA0007	80	100,00
ISSR_RPI_MFT_BCK	80	100,00
BSE_CLR	80	100,00
BCST_SR	80	100,00
VTBFHAPO	79	100,00
VDREPAYMENT	79	100,00
TRDT_SEC_FLOW	79	100,00
T552A	79	100,00
SOOD	79	100,00
PA0015	79	100,00
GRIX	79	100,00
GEOLOC	79	100,00
BSAK	79	100,00
ZZKV2O	78	100,00
TFAT	78	100,00
TEVEN	78	100,00
TCPSPPTL	78	100,00
LKSML	78	100,00
GRIX_VARP	78	100,00
GLT0	78	100,00
BSIK	78	100,00
BSAD	78	100,00
VIAK20	77	100,00
VIAK07	77	100,00
T706H	77	99,98
PA2007	77	100,00
PA0128	77	100,00
DFKKKO	77	100,00
DB6GSDTBS	77	100,00
CAT_VARS	77	100,00

TST01	76	100,00
TSP01	76	100,00
TRST_CLASFLO	76	100,00
TRQT_FLOW	76	100,00
PA0302	76	100,00
PA0000	76	100,00
IBROSTMS	76	100,00
HRSCONT	76	100,00
DFKKSUMC	76	100,00
DFKKRH	76	100,00
ANEP	76	100,00
T5D16	75	100,00
PA0020	75	100,00
DF40D	75	100,00
DB6CSTRACE	75	100,00
ZZTIME2	74	100,00
ZZISDEC	74	100,00
VWBEPI	74	100,00
T706V	74	100,00
T5UT3	74	100,00
SEOCLASSDF	74	100,00
PD3DSME	74	100,00
P01T_ADMIN_STAT	74	100,00
DD02L	74	100,00
TCVIEW	73	99,98
T706U	73	100,00
PPOIX	73	100,00
FKKMAZE	73	100,00
ECSCR_XML	73	100,00
DPAYP	73	100,00
DKKOP	73	100,00
BSIS	73	100,00
VDBEPI_EU	72	100,00
TSPEVJOB	72	100,00
T511	72	100,00
PD3FLAG	72	100,00
PA0019	72	100,00
EWUFI_SOP	72	100,00
DD01L	72	100,00
ZZISDEA	71	100,00
VDCFWAREHOUSE	71	100,00
VDBEPI	71	100,00
TRQT_BUSTRANS	71	100,00
T558A	71	100,00
DIMAPARSCPOS	71	100,00
DFKKZPT	71	100,00
CCSELTAB	71	100,00
TRQT_TRAFLOW	70	100,00
SOTR_TEXTU	70	100,00
PA0010	70	100,00
OCSPATNTCI	70	100,00
GRIX_SEL	70	100,00

EMMA_INT	70	100,00
EMMA_HDR	70	100,00
CEPRINT	70	100,00
CEFORMS	70	100,00
PCL2	69	100,00
OBJH	69	100,00
DM41S	69	100,00
DFKKRPL	69	100,00
DF54S	69	100,00
TA21L	68	100,00
T512D	68	100,00
SWDSMLINES	68	100,00
SWDSBINDEF	68	100,00
PA0022	68	100,00
FKKDIPOTMP	68	100,00
DFKKZS	68	100,00
DF53S	68	100,00
AGKO	68	100,00
TSKT4	67	100,00
TREELOG	67	100,00
TCNVF	67	100,00
SEOTYPEPLS	67	100,00
RTXTF	67	100,00
RF048	67	100,00
PUTTB_SHD	67	100,00
IBROSTMPSUMNO	67	100,00
DD04L	67	100,00
BDLSAIF	67	100,00
ADR2	67	100,00
TTREEI	66	100,00
SMMAIN	66	100,00
SEOCOMPODF	66	100,00
PD3DBME	66	100,00
ICNV31L	66	100,00
DPAYH	66	100,00
DM03S	66	100,00
DD40L	66	100,00
DD09L	66	100,00
VVKK_TRANSF_REF	65	100,00
TA22RSS	65	100,00
SXBNFNODE	65	100,00
SMSELKRIT	65	99,96
SMPARAM	65	100,00
SFRELN	65	100,00
RTXTL	65	99,88
FINI	65	96,98
EWUFZ	65	99,87
DIMAPARPLAN	65	100,00
DDYTF	65	100,00
DB6CSLITERAL	65	100,00
CUSAH	65	100,00
BUT021_FS	65	100,00

BSIP	65	100,00
BDLSADATA	65	100,00
AGR_TCODES	65	99,97
VARID	64	100,00
TST03	64	100,00
TMSBUFREQ	64	99,99
T5P1R	64	100,00
T52C5	64	100,00
SEOMETAREL	64	100,00
LOANREF_V	64	100,00
DD29L	64	100,00
D020S	64	100,00
CWBCIFIXED	64	100,00
AGR_HIER	64	100,00
ZFSCD_FAMK	63	100,00
TODIR	63	100,00
T512G	63	100,00
SWOTDV	63	100,00
STERM_LINK	63	100,00
SOTR_TEXT	63	100,00
SDOKMEP	63	99,98
DDPRS	63	99,81
VDPNS	62	100,00
T5BTM	62	99,98
T022D	62	99,95
SWDSTEXT	62	100,00
STXL	62	100,00
ROOSFIELD	62	99,99
PPOPX	62	100,00
P01T_LST1	62	100,00
JEST	62	99,99
ECSCR_LINE	62	99,99
E071KF	62	100,00
DD32S	62	100,00
DD25L	62	100,00
DB6PMHT HD	62	100,00
CWBNTCI	62	100,00
TQ29A	61	100,00
TA22RSF1	61	100,00
T51T3	61	99,98
T512P	61	99,76
PPDIX	61	100,00
DOKCR	61	100,00
DDNTT	61	100,00
DB6CONVL	61	99,98
SMENSAPNEW	60	100,00
SEOSUBCODF	60	100,00
SEOREDEF	60	100,00
DYNPLOAD	60	100,00
DFKKZV	60	100,00
DD03L	60	100,00
UST10S	59	99,99

TFDIR	59	100,00
SMENCUSNEW	59	100,00
SERPTREE	59	100,00
MSSSOURCE	59	100,00
SWOTLV	58	100,00
DSYST	58	100,00
CWBCIHEAD	58	99,81
AGR_1251	58	100,00
ZZSTAT1	57	100,00
USR13	57	99,90
SMFIAP	57	100,00
SMENSAP	57	99,99
SMENAKTNEW	57	100,00
LOANREF_A	57	99,98
EWUFZP	57	99,74
ENLFDIR	57	100,00
E071KFINI	57	100,00
AGR_1250	57	100,00
TBZ3R	56	99,99
T854T	56	99,49
PCALAC	56	100,00
FKK_UMBPOS	56	100,00
DSYSL	56	99,99
DFKKRAT	56	100,00
CWBNTHEAD	56	99,98
CWBNTCONT	56	99,99
CUS_ATRH	56	100,00
TVDIR	55	100,00
T5UTD	55	99,98
SWOTDQ	55	100,00
SMBATCH	55	99,99
ONR00	55	99,98
FUPARAREF	55	100,00
DSYAS	55	100,00
CUS_ACTOBJ	55	100,00
AQGTS	55	99,99
T52DB	54	99,44
SWOTLQ	54	100,00
DOKIL	54	100,00
TA22RSSY	53	100,00
T549B	53	99,77
SEOSUBCO	53	100,00
FUNCT	53	99,99
DD27S	53	100,00
DB6PMHT	53	100,00
CUS_ACTH	53	99,98
T800T	52	99,84
T52D9	52	99,19
DD26S	52	100,00
DD08T	52	99,99
DD07L	52	100,00
AGR_HIERT	52	100,00

WBCROSSGT	51	100,00
T5US1	51	99,77
FEBRE	51	100,00
DSYAT	51	99,99
CWBCICATTR	51	96,96
AGR_TEXTS	51	99,89
TTREES	50	99,20
TMDIR	50	99,97
T802T	50	99,15
T52CE_T	50	99,98
STERM_COMM	50	100,00
SEOCASSTX	50	99,96
D020T	50	99,99
T512L	49	99,80
SWW_CONTOB	45	100,00
CATX	44	88,87
BALDAT	2	19,10
ZZKV2C	0	0,00
ZZKV1C	0	0,00
ZZISDEP	0	0,00
BALHDR	0	0,00
SNAPT	-1	-1,00
DDNTF	-1	-1,00

Table 1 – Compression Candidates (with compression ratios)

The sum of these 310 tables make up about **132,5 GB**. After the compression the sum has shrunk to **39,7 GB** which is a **compression of 70%**.

The tables marked **red** are “bad compression candidates” and should not be compressed.

6.2 Largest Tables

The following table shows the TOP 20 tables ordered descending by size before compression:

TABNAME	PCTPAGESSAVED	SIZE (used-comp.) [KB]	SIZE (used-uncomp.) [KB]
DFKKOPK	83	3.392.768	19.957.459
VVSCPOS	81	2.936.096	15.453.137
DFKKOP	81	2.547.056	13.405.558
BALDAT	2	11.394.688	11.627.233
FKKDIPOTMP	68	3.097.328	9.679.150
DFKKRPL	69	2.901.488	9.359.639
VVSCITEM	80	1.719.232	8.596.160
DFKKKO	77	1.584.256	6.888.070
DPAYH	66	2.293.740	6.746.294
DPAYP	73	1.714.524	6.350.089
ZZISDEA	71	1.074.988	3.706.855
BSIS	73	835.224	3.093.422
ZZKV1T	85	433.356	2.889.040
ZZKV2T	86	346.336	2.473.829
ZZKV1P	85	369.956	2.466.373

BSAK	79	287.784	1.370.400
BALHDR	0	1.119.708	1.119.708
PPOIX	73	290.612	1.076.341
ZZKV2A	81	171.764	904.021
COEJ	85	116.972	779.813
	70	38.627.876	127.942.591

Table 2 – TOP 20 tables (sorted by uncompressed size)

These 20 tables make up 92% in size of all compressed tables. The average compression rate of these tables is **70%**.

Therefore it might be a good idea not to select all the candidates but only the big tables with a good compression rate.

The two tables BALHDR and BALDAT are big, but don't compress well. These tables should be excluded from compression. If we exclude them from Table 2 we get a compression ratio of **77%** on the remaining 18 tables.

6.3 CPU usage

From theory there is the expectation that the CPU usage is increased when you use a compressed system. Therefore we had a look at the average CPU usage before and after the compression. The result is shown in the following figure:




Figure 4 - CPU-Usage

The values were taken from transaction ST06. They represent the daily averages.

There was no CPU bottleneck in the system. The average CPU usage was slightly higher after the compression.

6.4 Response times

Another interesting fact is the behaviour of the response times in such a compressed system. To get a comparison data from transaction ST03N were taken. Mainly the task types DIALOG and BATCH were compared, because they are of highest interest in a productive SAP R/3 system for the daily work business.

The data basis is the week before (CW 46) and the week after the compression (CW 47). The following table lists the results for all transactions in average.

Task Type	Calender Week (CW)	# steps	Avg. resp. [ms]	Avg. CPU [ms]	Avg. DB [ms]
Dialog	CW 46	361.036	246,8	99,6	78,4
	CW 47	483.750	183,2	83,3	48,6
Diff.			63,6	16,3	29,8
% Diff.			(25,8%)	(16,4%)	(38,0%)
Batch	CW 46	44.850	10.135,9	2.722,8	3.076,9
	CW 47	43.106	7.708,9	3.331,0	2.374,3
Diff.			2.427,0	-608,2	702,6
% Diff.			(23,9%)	(-22,3%)	(22,8%)

Table 3 - Overall Response Time Comparison (ST03N)

As we see the workload concerning batch is nearly the same measured in terms of steps in the two comparing periods. For the task type dialog we can see a little increase in week 47 compared to week 46. There may be different reasons. Week 47 is the last week before month closing. So there might be some preparational work. And the better response times may have led to a bigger throughput of transactions, since the growth in the transaction steps is similar to the decrease in response time measured in percentage.

We see that for task type dialog both, avg. CPU and avg. DB times, which make up the most important parts of the total response time are decreasing while for batch processing the avg. CPU time is increasing for about 22%.

6.5 Dialog Response Times

To go a little more into detail, we will have a look at the dialog transactions and sort out the TOPS and FLOPS regarding the average response time.

Using this view we can see which transactions gained and which lost from the compression most.

To have a fair comparison some restrictions were put on the selection:

- # steps > 100 (only transactions that are used often)
- # steps : difference between the two weeks should not be greater than 100 %

Transaction	#Steps			avg. Resp.			avg CPU			avg DB		
	CW46	CW47	%Diff	CW46	CW47	%Diff	CW46	CW47	%Diff	CW46	CW47	%Diff
F.01	548	564	2,9	358,5	63,3	-82,3	93,3	19,9	-78,7	118,2	26,9	-77,2
FP05	748	534	-28,6	554,1	140,6	-74,6	42,3	25,7	-39,2	493,5	98,5	-80,0
FP30	533	631	18,4	750,0	191,3	-74,5	18,6	15,2	-18,3	712,7	167,2	-76,5
SQ01	631	596	-5,5	8.506,7	2.561,3	-69,9	2.405,5	778,1	-67,7	5.394,8	1.412,9	-73,8
KCLP	1.095	2.063	88,4	412,6	148,9	-63,9	56,7	48,9	-13,8	353,9	69,1	-80,5
F-43	3.163	4.136	30,8	304,1	110,2	-63,8	33,6	29,2	-13,1	23,6	14,7	-37,7
ZZSBH	2.392	3.854	61,1	1.975,9	787,3	-60,2	519,8	127,2	-75,5	847,1	307,5	-63,7
FBL1N	1.847	2.043	10,6	274,0	121,8	-55,5	47,6	44,5	-6,5	99,6	61,9	-37,9
PA30	11.713	21.062	79,8	234,2	107,7	-54,0	45,6	42,4	-7,0	49,5	27,6	-44,2
SE16	686	935	36,3	1.218,6	603,7	-50,5	87,3	59,9	-31,4	913,1	386,9	-57,6

Table 4 - TOPS dialog transactions with decreased response times

Table 4 shows that some transactions have dramatical gains in performance. All top 10 transactions reduce their response time more than 50%. There is a reduction in avg. CPU time and even more in avg. DB time.

For some of the transactions like SQ01 and SE16 the results might not be well comparable, because there might be a quite different workload behind.

Transaction	#Steps			avg. Resp.			avg CPU			avg DB		
	CW46	CW47	%Diff	CW46	CW47	%Diff	CW46	CW47	%Diff	CW46	CW47	%Diff
DB6SPACE	653	654	0,2	2.190,5	12.202,1	457,0	162,6	161,1	-0,9	316,4	98,3	-68,9
FO62	1.574	1.786	13,5	65,6	157,5	140,1	28,1	20,3	-27,8	24,2	10,2	-57,9
F110	3.228	2.170	-32,8	72,5	104,6	44,3	30,0	52,1	73,7	21,9	28,1	28,3
FP04	840	597	-28,9	150,1	162,7	8,4	32,9	32,8	-0,3	72,7	54,0	-25,7
F-02	5.947	5.673	-4,6	53,1	55,5	4,5	32,0	34,1	6,6	14,8	13,5	-8,8
INSOCHANGE	4.288	3.992	-6,9	346,9	360,2	3,8	165,5	157,0	-5,1	53,6	42,4	-20,9
PA20	2.000	2.665	33,3	290,7	297,1	2,2	54,9	55,3	0,7	50,1	30,6	-38,9
PC00_M01_CKTO	144.527	186.334	28,9	6,8	6,9	1,5	5,1	5,6	9,8	1,7	1,4	-17,6

Table 5 - FLOPS dialog transactions with increased response times

Table 5 shows the transactions that are the flops regarding the comparison of response times. But there are only 5 transactions that suffer quite a lot (more than 10%). If we have a closer look to the avg. CPU and avg. DB times we see that the avg. DB time is increasing only for transaction F110.

Another interesting question is if there is a dependency between the change of response time and the rate of DB changes a transaction is performing.

The following chart shows the data gathered in a graphical form.

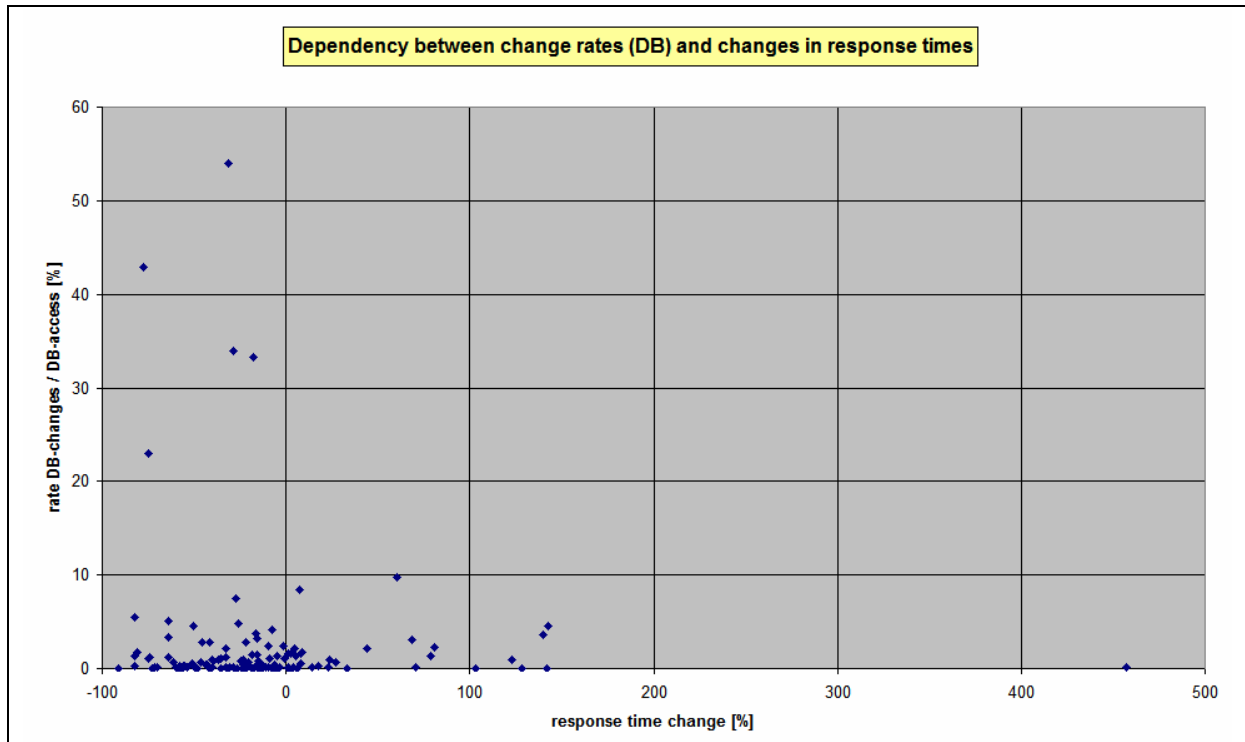


Figure 5 - Dependency between change rates (DB) and changes in response time

The data points show no strong correlation between the DB change rate and the increase or decrease of the response time. There are transactions that decreased in response time even when the DB change rate is high.

The whole picture looks distributed quite randomly. The only remarkable thing is that there are no points with coordinates (>0 , >10).

Since the DB-change rates comprise all types of changes (ins, upd, del), this might be a point where investigation should go deeper in future. From theory there is the expectation of a dependency between change of response times and amount of updates. Looking at all DB-changes this dependency can not be discovered for this data sample.

6.6 Batch Response Times

To get a quite fair comparison between the two weeks concerning batch jobs, there are some restrictions that were put on the comparison candidates. First it should be sure that the jobs were doing similar things. Furthermore it should be typical jobs, not special one.

The following table shows the reports that were taken for a week-to-week comparison:

Report	#Steps			avg. Resp.			avg CPU			avg DB		
	CW46	CW47	%Diff	CW46	CW47	%Diff	CW46	CW47	%Diff	CW46	CW47	%Diff
RFVIBL00	7	7	0,0	1.536,1	1.397,4	-9,0	54,3	57,1	5,2	463,9	401,6	-13,4
RKGALGA15	12	8	-33,3	116.578,7	82.655,3	-29,1	106.025,0	74.121,3	-30,1	8.849,0	7.574,3	-14,4
RKGALGA35	27	18	-33,3	46.387,1	56.976,9	22,8	35.083,7	40.641,1	15,8	5.448,5	4.704,2	-13,7
RPTCC106	1.788	1.623	-9,2	287,5	158,1	-45,0	58,7	55,7	-5,1	100,7	49,1	-51,2
RPTIME00	7	7	0,0	314.952,7	321.170,9	2,0	193.890,0	209.998,6	8,3	126.993,3	117.568,3	-7,4
RSAL_BATCH_TO OL_DISPATCHING	153	138	-9,8	18.825,0	17.277,7	-8,2	3.140,1	3.001,3	-4,4	7.971,9	6.000,0	-24,7
RSBTCRTE	4.747	7.451	57,0	73,7	52,7	-28,5	23,1	21,4	-7,4	31,5	24,6	-21,9
RSCOLL00	153	138	-9,8	12.043,1	21.535,9	78,8	2.827,3	3.791,3	34,1	2.354,2	6.128,4	160,3
RSCONN01	996	814	-18,3	115,6	97,5	-15,7	41,6	40,3	-3,1	47,6	43,9	-7,8
SAPCDT45	1.788	1.623	-9,2	281,2	270,9	-3,7	47,9	48,4	1,0	227,9	221,0	-3,0
ZFP00181	7	7	0,0	37.764,1	25.362,7	-32,8	4.994,3	4.561,4	-8,7	33.362,1	21.234,7	-36,4
ZFP00182	7	7	0,0	7.051,4	4.955,1	-29,7	2.152,9	2.170,0	0,8	5.174,7	3.150,4	-39,1
ZRP00071	7	7	0,0	1.313.859,1	1.442.696,0	9,8	1.295.620,0	1.410.561,4	8,9	23.653,0	36.740,9	55,3
ZRP00351	7	7	0,0	12.646,1	10.426,4	-17,6	5.024,3	4.901,4	-2,4	8.594,1	6.423,3	-25,3

Table 6 - Batch task comparison (sorted by report name)

As you can see most of the jobs have been better than before the compression, but three of them have suffered considerably regarding their response time.

RSCOLL00 is gathering statistical monitoring data into table MONI. This job has increased avg. DB time considerably.

7 Appendix

7.1 Definitions

To know how the numbers were calculated here are the formulas for the compression-factor and the compression-ratio

$$\text{C-factor (CF)} = \frac{\text{Bytes}(\text{uncomp.})}{\text{Bytes}(\text{comp.})} \quad \text{Eq. 1}$$

$$\text{Comp-Ratio (CR)} = \left(1 - \frac{\text{Bytes}(\text{comp.})}{\text{Bytes}(\text{uncomp.})}\right) * 100 \quad [\%] \quad \text{Eq. 2}$$

and with concern of Eq. 1 we get CR as a function of CF:

$$\text{CR} = f(\text{CF}) = \left(1 - \frac{1}{\text{CF}}\right) * 100 \quad \text{Eq. 3}$$

and by transforming Eq. 3 we get CF as a function of CR as:

$$\text{CF} = f(\text{CR}) = \frac{100}{100 - \text{CR}} \quad \text{Eq. 4}$$

7.2 DB2 Registry Variables

```
[e] DB2BQTIME=1
[e] DB2BQTRY=60
[e] DB2CHKPTR=OFF
[e] DB2CODEPAGE=819
[e] DB2COMM=TCPIP
[e] DB2DBDFT=PR1
[e] DB2IQTIME=5
[e] DB2RQTIME=5
[i] DB2_OPT_MAX_TEMP_SIZE=10240 [DB2_WORKLOAD]
[i] DB2_WORKLOAD=SAP
[i] DB2_TRUNCATE_REUSESTORAGE=IMPORT [DB2_WORKLOAD]
[i] DB2_MDC_ROLLOUT=YES [DB2_WORKLOAD]
[i] DB2_SKIPINSERTED=YES [DB2_WORKLOAD]
[i] DB2_VIEW_REOPT_VALUES=YES [DB2_WORKLOAD]
[i] DB2_OBJECT_TABLE_ENTRIES=65532 [DB2_WORKLOAD]
[i] DB2_OPTPROFILE=YES [DB2_WORKLOAD]
[i] DB2_IMPLICIT_UNICODE=YES [DB2_WORKLOAD]
[i] DB2_INLIST_TO_NLJN=YES [DB2_WORKLOAD]
[i] DB2_MINIMIZE_LISTPREFETCH=YES [DB2_WORKLOAD]
[i] DB2_UPDATE_PART_KEY=YES [DB2_WORKLOAD]
[i] DB2_REDUCED_OPTIMIZATION=4, INDEX, JOIN, NO_TQ_FACT, NO_HSJN_BUILD_FACT,
STARJN_CARD_SKEW, NO_SORT_MGJOIN [DB2_WORKLOAD]
[i] DB2NOTIFYVERBOSE=YES [DB2_WORKLOAD]
[i] DB2_INTERESTING_KEYS=YES [DB2_WORKLOAD]
[i] DB2_EVALUNCOMMITTED=YES [DB2_WORKLOAD]
[i] DB2_DISABLE_FLUSH_LOG=ON
[i] DB2_ANTIJOIN=EXTEND [DB2_WORKLOAD]
```



```

[i] DB2_STRIPED_CONTAINERS=ON
[i] DB2_HASH_JOIN=ON
[i] DB2MEMMAXFREE=2000000 [DB2_WORKLOAD]
[i] DB2ENVLIST=Not a terminal INSTHOME SAPSYSTEMNAME dbs_db6_schema DIR_LIBRARY LIBPATH
[i] DB2_RR_TO_RS=YES [DB2_WORKLOAD]
[i] DB2_BLOCK_ON_LOG_DISK_FULL=ON
[i] DB2_FORCE_FCM_BP=YES [DB2_WORKLOAD]
[i] DB2COMM=TCPIP [0]
[g] DB2SYSTEM=sapprod
[g] DB2INSTDEF=kwsoft
[g] DB2ADMINSERVER=dasusr1

```

7.3 DBM Configuration

Database Manager Configuration

Node type = Enterprise Server Edition with local and remote clients

```

Database manager configuration release level          = 0x0b00

CPU speed (millisec/instruction)                    (CPUSPEED) = 5.235149e-07
Communications bandwidth (MB/sec)                   (COMM_BANDWIDTH) = 2.000000e+00

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/PR1/sqlllib/java/jdk64

Diagnostic error capture level                       (DIAGLEVEL) = 3
Notify Level                                         (NOTIFYLEVEL) = 3
Diagnostic data directory path                       (DIAGPATH) = /db2/PR1/sqlllib/db2dumpp

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) = DBPR1ADM
SYSCTRL group name                                  (SYSCTRL_GROUP) = DBPR1CTL
SYSMAINT group name                                 (SYSMAINT_GROUP) = DBPR1MNT
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/PR1

Database monitor heap size (4KB)                    (MON_HEAP_SZ) = 256
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) = 0
Directory cache support                   (DIR_CACHE) = YES
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) = sapdb2PR1
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

```

7.4 DB Configuration

Database Configuration for Database PR1

```

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

Database territory                       = en_US
Database code page                       = 819
Database code set                        = ISO8859-1
Database country/region code             = 1
Database collating sequence              = IDENTITY
Alternate collating sequence              (ALT_COLLATE) =
Database page size                       = 4096

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 = NO

Log retain for recovery status = RECOVERY
User exit for logging status = YES

Self tuning memory (SELF_TUNING_MEM) = ON
Size of database shared memory (4KB) (DATABASE_MEMORY) = 1000000
Database memory threshold (DB_MEM_THRESH) = 10
Max storage for lock list (4KB) (LOCKLIST) = AUTOMATIC
Percent. of lock lists per application (MAXLOCKS) = AUTOMATIC
Package cache size (4KB) (PCKCACHESZ) = AUTOMATIC
Sort heap thres for shared sorts (4KB) (SHEAPTHRES_SHR) = AUTOMATIC
Sort list heap (4KB) (SORTHEAP) = AUTOMATIC

Database heap (4KB) (DBHEAP) = 65000
Catalog cache size (4KB) (CATALOGCACHE_SZ) = 2560
Log buffer size (4KB) (LOGBUFSZ) = 1024
Utilities heap size (4KB) (UTIL_HEAP_SZ) = 1000
Buffer pool size (pages) (BUFFPAGE) = 90000
Max size of appl. group mem set (4KB) (APPGROUP_MEM_SZ) = 40000
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) = 60000
Default application heap (4KB) (APPLHEAPSZ) = 3072
Statistics heap size (4KB) (STAT_HEAP_SZ) = 15000

Interval for checking deadlock (ms) (DLCHKTIME) = 300000
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) = AUTOMATIC

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) = 65520
Number of primary log files (LOGPRIMARY) = 19
Number of secondary log files (LOGSECOND) = 81
Changed path to log files (NEWLOGPATH) =
Path to log files =
/db2/PR1/db2pr1/NODE0000/SQL00001/SQLOGDIR/
Overflow log path (OVERFLOWLOGPATH) =
Mirror log path (MIRRORLOGPATH) =
First active log file = S0012360.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 ckcpt (SOFTMAX) = 300

```

```

Log retain for recovery enabled          (LOGRETAIN) = RECOVERY
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) = TSM:MCSAP3
Options for logarchmeth1                 (LOGARCHOPT1) =
Second log archive method                (LOGARCHMETH2) = OFF
Options for logarchmeth2                 (LOGARCHOPT2) =
Failover log archive path                (FAILARCHPATH) = /db2/PR1/log_fail/
Number of log archive retries on error   (NUMARCHRETRY) = 5
Log archive retry Delay (secs)           (ARCHRETRYDELAY) = 20
Vendor options                           (VENDOROPT) =

Auto restart enabled                     (AUORESTART) = ON
Index re-creation time and redo index build (INDEXREC) = 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) = ON
  Automatic database backup               (AUTO_DB_BACKUP) = OFF
  Automatic table maintenance            (AUTO_TBL_MAINT) = ON
  Automatic runstats                     (AUTO_RUNSTATS) = ON
  Automatic statistics profiling          (AUTO_STATS_PROF) = OFF
  Automatic profile updates               (AUTO_PROF_UPD) = OFF
  Automatic reorganization                (AUTO_REORG) = OFF

```

8 List of Literature

- [1] Ziv, Jacob; Lempel, Abraham: A Universal Algorithm for Sequential Data Compression, IEEE Transactions on Information Theory, Vol. IT-23, No. 3, p. 337-343, May 1977
http://www.cs.duke.edu/courses/spring03/cps296.5/papers/ziv_lempel_1977_universal_algorithm.pdf
- [2] IBM DB2 Laboratory, Toronto: Compression in DB2 Viper, May 2006
- [3] SAP: Database Upgrade Guide, Migration to Version 9 of IBM DB2 for Linux, UNIX, and Windows; Document Version 1.00 – Aug 11th, 2006
(find in the SAP Service Marketplace, Quick-Link: instguides, > Other Documentation > Database Upgrades > DB2 UDB)
- [4] Heinrich, Johannes (SAP AG): New Features in DB2 UDB V9 – Part 2
<https://www.sdn.sap.com/irj/sdn/weblogs?blog=/pub/wlg/3750>
- [5] Chen, Whei-Jen; Donner, Jochen; König, Edgardo G.; Konno, Masako; Tang, Beck; Wang, Xiaomei: SAP Solutions on IBM DB2 UDB V8.2.2 - Handbook, IBM Redbook SG24-6765-00
<http://w3.itso.ibm.com/abstracts/sg246765.html?Open>

SAP Notes (<http://service.sap.com/notes>)

[Note 980067 - DB6: Using DB2 9 Row Compression](#)

[Note 930487 - DB6: Using DB2 9 with SAP Software](#)

[Note 905614 - DB6: R3load -loadprocedure fast COMPRESS](#)

[Note 817709 - DB6: DB6CONV: Complete version history](#)

[Note 486559 - DB2DART: Reducing high water mark](#)

[Note 362325 - DB6: Table conversion using DB6CONV](#)

[Note 152531 - High water mark](#)

9 Table of Figures / Tables / Equations

9.1 Figures

FIGURE 1 – COMPONENTS OF THE SAP SYSTEM..... 7
 FIGURE 2 - PRODUCTIVE SYSTEM ENVIRONMENT 7
 FIGURE 3 - COMPRESSION ALGORITHM 8
 FIGURE 4 - CPU-USAGE..... 19
 FIGURE 5 - DEPENDENCY BETWEEN CHANGE RATES (DB) AND CHANGES IN RESPONSE TIME 22

9.2 Tables

TABLE 1 – COMPRESSION CANDIDATES (WITH COMPRESSION RATIOS) 18
 TABLE 2 – TOP 20 TABLES (SORTED BY UNCOMPRESSED SIZE)..... 19
 TABLE 3 - OVERALL RESPONSE TIME COMPARISON (ST03N) 20
 TABLE 4 - TOPS DIALOG TRANSACTIONS WITH DECREASED RESPONSE TIMES..... 21
 TABLE 5 - FLOPS DIALOG TRANSACTIONS WITH INCREASED RESPONSE TIMES 21
 TABLE 6 - BATCH TASK COMPARISON (SORTED BY REPORT NAME)..... 23

9.3 Equations

$$C\text{-FACTOR (CF)} = \frac{\text{Bytes}(\text{uncomp.})}{\text{Bytes}(\text{comp.})} \quad \text{EQ. 1..... 24}$$

$$\text{COMP-RATIO (CR)} = \left(1 - \frac{\text{Bytes}(\text{comp.})}{\text{Bytes}(\text{uncomp.})}\right) * 100 \quad [\%] \quad \text{EQ. 2..... 24}$$

$$\text{CR} = F(\text{CF}) = \left(1 - \frac{1}{\text{CF}}\right) * 100 \quad \text{EQ. 3..... 24}$$

$$\text{CF} = F(\text{CR}) = \frac{100}{100 - \text{CR}} \quad \text{EQ. 4..... 24}$$