

Support Workshop
BW Feature Pack

SAP AG
Berlin, June 2007

THE BEST-RUN BUSINESSES RUN SAP™



New DB Features for BI

BI / Netweaver Integration

Support Issues



BW Features in MaxDB 7.6. :

Table Clustering

- Logical Table Clustering
- Physical Table Clustering

Table Compression

Join Optimization

The BW Feature Pack is available for BI systems as of MaxDB Version 7.6 Support Package 01. This BW Feature Pack improves the performance of BI systems.

Logical Cluster (see next slide)

Physical table cluster

The CLUSTER table attribute is included in the tables. As a result of this attribute, the table contents are no longer distributed evenly among all volumes in individual blocks of eight kilobyte pages. Instead, if possible, they are stored among the volumes in larger connected blocks.

Table Compression

The PACKED table attribute causes integer values to be compressed and stored and as a result, they require less memory space on the volumes. The results in a faster I/O access and reduced displacement in the data cache.

Logical Table Clustering



- Use CLUSTER columns internally as primary key
- Ensure uniqueness of the primary key by adding a postfix sequence
- Only for tables with no user defined primary key

```
CREATE TABLE clu
col1    VARCHAR(20),
col2    VARCHAR(20),
created VARCHAR(8 ))
CLUSTER (created)
```

```
1: (pos 00081) key(13): 20323030 36303130 31000000 01 ' 20060101....'
2: (pos 00165) key(13): 20323030 36303130 31000000 02 ' 20060101....'
3: (pos 00249) key(13): 20323030 36303130 31000000 03 ' 20060101....'
4: (pos 00277) key(13): 20323030 36303130 31000000 04 ' 20060101....'
5: (pos 00305) key(13): 20323030 36303130 31000000 05 ' 20060101....'
6: (pos 00333) key(13): 20323030 36303130 31000000 06 ' 20060101....'
7: (pos 00361) key(13): 20323030 36303130 31000000 07 ' 20060101....'
8: (pos 00389) key(13): 20323030 36303130 31000000 08 ' 20060101....'
9: (pos 00417) key(13): 20323030 36303130 31000000 09 ' 20060101....'
...
18: (pos 00109) key(13): 20323030 36303130 32000000 01 ' 20060102....'
19: (pos 00193) key(13): 20323030 36303130 32000000 02 ' 20060102....'
...
20: (pos 00137) key(13): 20323030 36303130 33000000 01 ' 20060103....'
```

© SAP 2007 / SWS 2007 BW Feature Pack / Page 4

A virtual key allows you to create a key for non-unique columns. From a technical point of view, this virtual key consists of a non-unique key field and a unique sequential number (SEQUENCE). If a table has a virtual key, logical clustering of the table also takes place.

This type of virtual key is created in the BW environment in connection with the physical clustering of the fact tables.

The report from Note 983845 adds the CLUSTER attribute to the fact tables and creates the virtual key.

Benefit: Logical Table Clustering

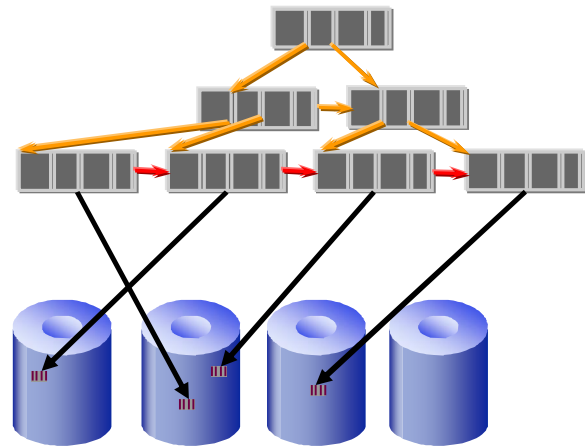
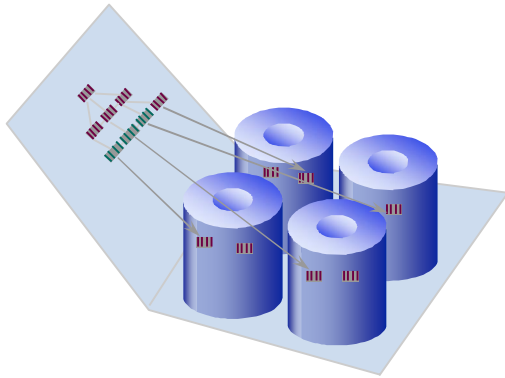


- Join via primary key instead of join via index
- Field experience: 50% faster join (precondition: without IO)

Challenge: Physical Table Clustering



- Database tasks read blocks from disks in 8KB units
- Scans don't benefit from larger block sizes of storage systems and read-ahead options




Physical Table Clustering (Write)




- Sort pages in cache via the right page pointer and write them down in clusters
- Write big clusters to the end of the volume. They fill the volume from the end to the beginning. 8 KB clusters fill the volumes from start to the end

Data Cache



Converter

Page	Volume	Offset
4711	1	9857
...		
7598	1	9858
...		
9328	1	9859
...		
9587	1	9860




Physical Table Clustering (Read)




- Read page numbers from separators in index level (1) and check the block positions in converter
- Read pages of one cluster with one I/O

Data Cache



Converter

Page	Volume	Offset
4711	1	9857
...		
7598	1	9858
...		
9328	1	9859
...		
9587	1	9860



Physical Table Clustering (System Views)



- The view **DATAVOLUMES** shows the **CLUSTERAREASIZE**
- The view **FILES** shows the **CLUSTER** attribute
- The views **TABLESTORAGEDETAILS** and **INDEXSTORAGEDETAILS** show the number of clusters in the column **CLUSTERCOUNT** (Select only the columns **SCHEMANAME**, **TABLERNAME**, **INDEXNAME** and **CLUSTERCOUNT** to prevent a full table scan)
- Optimum number of clusters would be:
$$\frac{\langle \text{TREELEAVESIZE} \rangle}{8} / \langle \text{DATA_IO_BLOCK_COUNT} \rangle$$

e.g. $260448 / 8 / 64 \sim 509$

```
select d.tablename, '' as clustercount,
       f.type, f.entrycount, f.treeleavessize, f.treeindexsize, f.lobsize, f.clustered
from tablestagedetails d, files f, tables t
where d.schemaname = t.schemaname
and d.tablename = t.tablename
and f.fileid = t.tableid
and t.tablename = 'MYCLUSTER'
union all
select d.tablename, d.indexname, d.clustercount,
       f.type, f.entrycount, f.treeleavessize, f.treeindexsize, f.lobsize, f.clustered
from indexstagedetails d, files f, tables t
where d.schemaname = t.schemaname
and d.tablename = t.tablename
and f.primaryfileid = t.tableid
and t.tablename = 'MYCLUSTER'
```

TABLERNAME	EXPRESSION1	CLUSTERCOUNT	TYPE	ENTRYCOUNT	TREELEAVESIZE	TREEINDEXSIZE	LOBSIZE	CLUSTERED
MYCLUSTER		511	TABLE	2500000	260448	912	0	YES
MYCLUSTER	MYCLUSTER_I	1814	INDEX	230172	80528	7368	?	YES

© SAP 2007 / SWS 2007 BW Feature Pack / Page 9

You can check whether a table has the **CLUSTER** attribute by looking at the **CLUSTERED** column in the **FILES** system table.

You use selects for tables **TABLESTORAGEDETAILS** or **INDEXSTORAGEDETAILS** to evaluate the cluster quality of a table. Note that SQL accesses to these system tables where no qualification is specified in the **WHERE** condition for table or index names or one of the columns can lead to long runtimes.

You can use the following select to quickly determine the cluster properties for tables:
select **SCHEMANAME**, **TABLERNAME**, **CLUSTERCOUNT** from **TABLESTORAGEDETAILS**
where **SCHEMANAME** = <schemaname> and **TABLERNAME** = <tablename>

Compared with **TREELEAVESIZE**:
select **TREELEAVESIZE** from **FILES**, **ROOTS**
where **FILES.FILEID** = **ROOTS.TABLEID**
and **ROOTS.TABLERNAME** = <tablename>
and **ROOTS.SCHEMANAME** = <schemaname>

The **CLUSTERCOUNT** value counts the data block areas that are connected and sorted in ascending order; however, these data block areas do not necessarily have to be located in the clustered area of a volume.

TREELEAVESIZE (in kilobytes) is divided by the page size (in kilobytes), which is divided by **CLUSTERCOUNT**, and this determines the cluster quality. The closer the result is to the **DATA_IO_BLOCK_COUNT**, the better the clustering. A result with a value of 1 shows that the table is not clustered and stored.

Benefit: Physical Table Clustering



- Database tasks read blocks from disks in cluster units
- Scans benefit from larger block sizes of storage systems and read-ahead options
- Field experience: 5-6 times in average (20 times) faster scan





Table Compression: numeric columns



- Compression of numeric columns
 - In general fact tables have only numeric columns
 - If fact table has no key figures with data type FLOAT:
Length indicator occupies 4 Bit instead of 1 Byte, because maximum value length is 15
 - Two digits are stored in 7 Bits instead of 1 Byte

Non Packed:



Packed:



© SAP 2007 / SWS 2007 BW Feature Pack / Page 11

You can check whether a table has the PACKED attribute and therefore, whether integer values are packed and stored by looking at the CLUSTERED column in the FILES system table.

Table Compression: 0-values and NULL



- Additional compression of frequent 0-values
 - Length occupies 4 Bit instead of 1 Byte

Non Packed:

length 1	0
----------	---

Packed:

0

- Omit Defined Bytes for NULL values
 - Non packed fields occupy a Defined Byte to identify NULL values
 - This Defined Byte is not needed if the database knows that records will not have NULL values in the column
 - Works only with NOT NULL

Table Compression: Requirements + Benefit




- Requirements
 - All non primary key columns are numeric
 - NOT NULL
 - Maximum 15 Byte per column
 - PACKED attribute will be ignored if requirements are not fulfilled

- Benefit:
 - Field Experience: 60-70% smaller tables and 30-40% smaller indexes



- Joins via hash tables make sense, if the ratio between the size of the intermediate join result and the size of the next join table exceeds a certain value.
- The parameter `OPTIMIZE_JOIN_HASH_MINIMAL_RATIO` defines the minimal ratio to use hashing for the next table.
- New: The decision if a hash join should be executed takes into account the intermediate result size instead of the last table.
- You can override this parameter with the hint `HASHRATIO(<unsigned int>)`

```
SELECT /* HASHRATIO(1) */ ...  
FROM tab1, tab2  
WHERE tab1.key = tab2.field
```



- Benefit: 4-5 Seconds for each 1 Mio rows in result sets.
Field Experience: select 2 times faster

Move Clouds of Dimensions in Optimizer Strategy



- Use knowledge about fact tables

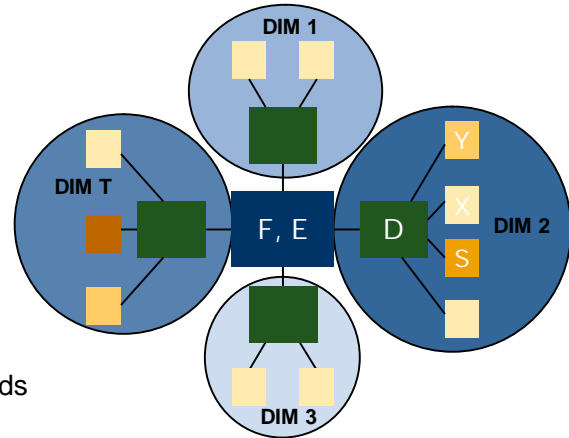
The MaxDB optimizer is able to identify a fact table:

```
ALTER TABLE <tablename> FACT | DIMENSION
```

- Join from first dimension cloud to fact table

The join between two dimension clouds can significantly increase the temporary result to be joined with the fact table.

The system now ensures the join between one dimension cloud to the fact table before joining the other dimension clouds.



- Benefit: Prevent from long runtimes of joins due to undesirable joins between dimension clouds

New DB Features for BI

BI / Netweaver Integration

Support Issues



- Requirements
 - MaxDB 7.6.01
 - ABAP corrections for BW 3.x and BI 7.x, see CSN note 983845
 - Available in Support Packages (Basis and BW)
 - in urgent cases use the correction instruction
- New cube and aggregate tables (fact tables) will be created automatically with BW Feature Pack attributes (clustered, packed)
- Previous fact tables have to be converted with the report RSDU_CLUSTER_FACT_ADA

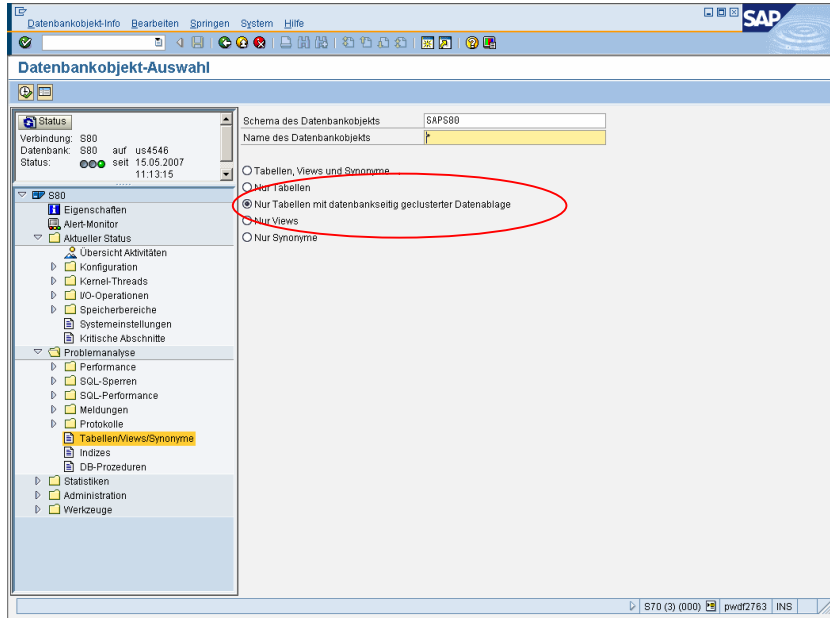
To use the BW Feature Pack, you must upgrade to MaxDB Version 7.6.01 or higher. For more information, see Note 983845. We recommend that you use the Save Data and Restore Data functions for the database before converting the fact tables.

If you use the Restore Data function, the tables are evenly distributed among the volumes. Fact table clustering can then be optimized.

Monitoring of packed and clustered tables



- Transaction DB50 -> Problem Analysis -> Tables/Views/Synonyms



Cluster attribute in table definition



- Only the cluster key field is visible, not the sequence field (virtual key)

The screenshot shows the SAP Table Definition tool interface. The main window displays the table definition for 'SAPS80/BIC/FISQCAUF'. A message at the top states: 'Tabelle hat einen virtuellen Primärschlüssel. KEY-Spalten können daher Mehrfacheinträge haben.' (Table has a virtual primary key. KEY columns can therefore have multiple entries). The table definition table is as follows:

Spaltenname	Typ	Datentyp	Codebtp	Län.	D.	Zug.	Default	P.	Sc.	Erzeugungsd.	Zeit	Änder.
KEY_ISQCAUFP	OPT	NUMBER		10	0	SE...	0		1	11.11.2005	12:04:59	11.11.
KEY_ISQCAUFU	KEY	NUMBER		10	0	SE...	0		2	11.11.2005	12:04:59	11.11.
KEY_ISQCAUF1	OPT	NUMBER		10	0	SE...	0		3	11.11.2005	12:04:59	11.11.
KEY_ISQCAUF2	OPT	NUMBER		10	0	SE...	0		4	11.11.2005	12:04:59	11.11.
KEY_ISQCAUF3	OPT	NUMBER		10	0	SE...	0		5	11.11.2005	12:04:59	11.11.
KEY_ISQCAUF4	OPT	NUMBER		10	0	SE...	0		6	11.11.2005	12:04:59	11.11.
BBP_ASPGOU	OPT	NUMBER		17	3	SE...	0.000		8	11.11.2005	12:04:59	11.11.
BBP_POCT	OPT	NUMBER		17	3	SE...	0.000		9	11.11.2005	12:04:59	11.11.

Storage details of clustered table



- Factor: value corresponds with the benefit for scan

The screenshot shows the SAP 'Tabellen/View-Information' window. The left sidebar contains a tree view with 'Tabellen/Views/Synonyme' selected. The main area displays a table with the following columns: 'Tabelle / Indizes', 'Geclustert', 'Seiten (orakt)', 'Cluster', 'Faktor', 'Gepackte nurn', and 'Spalten'. The data row shows: 'SAP550/BIC/FISOCAUF', 'Ja', '5.981', '1.629', '3,67', and 'Ja'. The 'Faktor' and 'Gepackte nurn' columns are circled in red.

Tabelle / Indizes	Geclustert	Seiten (orakt)	Cluster	Faktor	Gepackte nurn	Spalten
SAP550/BIC/FISOCAUF	Ja	5.981	1.629	3,67		Ja

New DB Features for BI

BI / Netweaver Integration

Support Issues

Recommended MaxDB Version



- Min. Version to activate BW Feature Pack: 7.6.01
- Recommended min. version: 7.6.02.x, x >= 10
 - Miscellaneous error corrections for Hash Join, Parallel Join
 - Hash Join on Unique Index
 - Packed table with columns of datatype FLOAT



- Remember: page clusters are assembled during save point
- Updates and inserts inside the cluster split the cluster
- Short page clusters get lost after RESTORE DATA
- No automatic and dynamic reclustering at the moment (planned for 7.6.04)

Execution Plan of a BW Query



Only 1 dimension cloud before fact table

Join transition to fact table via clustered key

```
explain
SELECT
Tab...
```

TABNAME	COLUMN_OR_INDEX	STRATEGY	PAGECOUNT
S1	/BIO/SFISCPER~Z01 FISCPER	RANGE CONDITION FOR INDEX (USED INDEX COLUMN)	20
DT	/BIC/DSHD_ERGT...	JOIN VIA RANGE OF MULTIP... TABLE HASHED	1
F	SID_OFISCPER KEY_SHD_ERGT	(USED INDEX COLUMN) JOIN VIA KEY RANGE	16566
DU	DIMID	JOIN VIA KEY COLUMN TABLE HASHED	1
D5	DIMID	JOIN VIA KEY COLUMN TABLE HASHED	30
D4	DIMID	JOIN VIA KEY COLUMN TABLE HASHED	307
D2	DIMID	JOIN VIA KEY COLUMN TABLE HASHED	28
DP	DIMID	JOIN VIA KEY COLUMN TABLE HASHED	2
D1	DIMID	JOIN VIA KEY COLUMN TABLE HASHED	1
JDBC_CURSC		RESULT IS COPIED , CO...	112183

All tables behind the fact table are hashed



- Only EXPLAIN SEQUENCE provides information about Star Schema Search

```

explain sequence
SELECT
/*+
SHORT_SUM_VALUES
*/
"DU"."SID_OCCURRENCY" AS "S____009", "DU"."SID_OCOPASLQTU" AS "S____013",

```

TEXT			
6 SINGLE KEY	0.100E+01	0.104347E+0...	
7 SINGLE KEY	0.100E+01	0.101754E+0...	
8 SINGLE KEY	0.100E+01	0.101754E+0...	
9 SINGLE KEY	0.100E+01	0.203508E+0...	
JOIN SEQUENCE SEARCH = STAR SCHEMA			
SEARCH DIMENSION PLANS			
< COSTVALUE : 0.188790E+01 [2 1]			
SEARCH COMPLETE PLAN			
< COSTVALUE : 0.332521E+06 [3 9 5 6 7 ...			
< COSTVALUE : 0.113705E+06 [9 3 5 6 7 ...			
> COSTVALUE : 0.140621E+06 [5 3 9 6 7 ...			
\> COSTVALUE : 0.148614E+06 [6 3 9 5 7			

- Parameter OPTIMIZE_STAR_JOIN = YES
- Fact table has FACT attribut
 - FACT attribut get lost by CREATE TABLE ... LIKE
 - Set FACT attribut with report RSDU_SET_FACT_ATTR_ADA
- Only MaxDB 7.6.1: fact table has to be joined directly only with dimension tables (this precondition is dropped in MaxDB 7.6.2)

Thank you!





No part of this publication may be reproduced or transmitted in any form or for any purpose without the express permission of SAP AG. The information contained herein may be changed without prior notice.

Some software products marketed by SAP AG and its distributors contain proprietary software components of other software vendors.

SAP, R/3, mySAP, mySAP.com, xApps, xApp, SAP NetWeaver, Duet, Business ByDesign, ByDesign, PartnerEdge and other SAP products and services mentioned herein as well as their respective logos are trademarks or registered trademarks of SAP AG in Germany and in several other countries all over the world. All other product and service names mentioned and associated logos displayed are the trademarks of their respective companies. Data contained in this document serves informational purposes only. National product specifications may vary.

The information in this document is proprietary to SAP. This document is a preliminary version and not subject to your license agreement or any other agreement with SAP. This document contains only intended strategies, developments, and functionalities of the SAP® product and is not intended to be binding upon SAP to any particular course of business, product strategy, and/or development. SAP assumes no responsibility for errors or omissions in this document. SAP does not warrant the accuracy or completeness of the information, text, graphics, links, or other items contained within this material. This document is provided without a warranty of any kind, either express or implied, including but not limited to the implied warranties of merchantability, fitness for a particular purpose, or non-infringement.

SAP shall have no liability for damages of any kind including without limitation direct, special, indirect, or consequential damages that may result from the use of these materials. This limitation shall not apply in cases of intent or gross negligence.

The statutory liability for personal injury and defective products is not affected. SAP has no control over the information that you may access through the use of hot links contained in these materials and does not endorse your use of third-party Web pages nor provide any warranty whatsoever relating to third-party Web pages

Weitergabe und Vervielfältigung dieser Publikation oder von Teilen daraus sind, zu welchem Zweck und in welcher Form auch immer, ohne die ausdrückliche schriftliche Genehmigung durch SAP AG nicht gestattet. In dieser Publikation enthaltene Informationen können ohne vorherige Ankündigung geändert werden.

Einige von der SAP AG und deren Vertriebspartnern vertriebene Softwareprodukte können Softwarekomponenten umfassen, die Eigentum anderer Softwarehersteller sind.

SAP, R/3, mySAP, mySAP.com, xApps, xApp, SAP NetWeaver, Duet, Business ByDesign, ByDesign, PartnerEdge und andere in diesem Dokument erwähnte SAP-Produkte und Services sowie die dazugehörigen Logos sind Marken oder eingetragene Marken der SAP AG in Deutschland und in mehreren anderen Ländern weltweit. Alle anderen in diesem Dokument erwähnten Namen von Produkten und Services sowie die damit verbundenen Firmenlogos sind Marken der jeweiligen Unternehmen. Die Angaben im Text sind unverbindlich und dienen lediglich zu Informationszwecken. Produkte können länderspezifische Unterschiede aufweisen.

Die in diesem Dokument enthaltenen Informationen sind Eigentum von SAP. Dieses Dokument ist eine Vorabversion und unterliegt nicht Ihrer Lizenzvereinbarung oder einer anderen Vereinbarung mit SAP. Dieses Dokument enthält nur vorgesehene Strategien, Entwicklungen und Funktionen des SAP®-Produkts und ist für SAP nicht bindend, einen bestimmten Geschäftsweg, eine Produktstrategie bzw. -entwicklung einzuschlagen. SAP übernimmt keine Verantwortung für Fehler oder Auslassungen in diesen Materialien. SAP garantiert nicht die Richtigkeit oder Vollständigkeit der Informationen, Texte, Grafiken, Links oder anderer in diesen Materialien enthaltenen Elemente. Diese Publikation wird ohne jegliche Gewähr, weder ausdrücklich noch stillschweigend, bereitgestellt. Dies gilt u. a., aber nicht ausschließlich, hinsichtlich der Gewährleistung der Marktgängigkeit und der Eignung für einen bestimmten Zweck sowie für die Gewährleistung der Nichtverletzung geltenden Rechts.

SAP übernimmt keine Haftung für Schäden jeglicher Art, einschließlich und ohne Einschränkung für direkte, spezielle, indirekte oder Folgeschäden im Zusammenhang mit der Verwendung dieser Unterlagen. Diese Einschränkung gilt nicht bei Vorsatz oder grober Fahrlässigkeit.

Die gesetzliche Haftung bei Personenschäden oder die Produkthaftung bleibt unberührt. Die Informationen, auf die Sie möglicherweise über die in diesem Material enthaltenen Hotlinks zugreifen, unterliegen nicht dem Einfluss von SAP, und SAP unterstützt nicht die Nutzung von Internetseiten Dritter durch Sie und gibt keinerlei Gewährleistungen oder Zusagen über Internetseiten Dritter ab.

Alle Rechte vorbehalten.