SAP® MaxDB™
Optimized for SAP Business Warehouse

Expert Session 9

MaxDB/liveCache Development Support
Christiane Hienger
Melanie Handreck
May 4, 2010

THE BEST-RUN BUSINESSES RUN SAP™

Agenda

1. Introduction
2. MaxDB Data Storage for BW
   - Physical Table Clustering
   - Table Compression
3. Logical Table Clustering
4. MaxDB Optimizer
   - Optimized Star-Join-Processing
   - Optimized Hash Join
5. Performance Analysis and Monitoring
6. Summary

© 2010 SAP AG. All rights reserved. / Page 2

The target group of this session are customers with SAP Business Warehouse systems running on SAP MaxDB, SAP partners and SAP employees who are doing BW performance analysis on SAP MaxDB systems.

As of SAP MaxDB version 7.6, MaxDB offers special features, which significantly improve the performance of SAP Business Warehouse applications.

This session provides information about these MaxDB features called **SAP MaxDB BW Feature Pack.**

The SAP Business Warehouse application itself is not part of this session.

This session is based on the internal development System S70. It is a SAP NW 7.0 system with limited BW functionality. The MaxDB Version is 7.7.07.

This session consists of a theoretical part and a second part where a performance analysis including checks of the usage of the BW Feature Pack is done directly on the system S70.

Agenda

1. Introduction
2. MaxDB Data Storage for BW
   - Physical Table Clustering
   - Table Compression
3. Logical Table Clustering
4. MaxDB Optimizer
   - Optimized Star-Join-Processing
   - Optimized Hash Join
5. Performance Analysis and Monitoring
6. Summary

© 2010 SAP AG. All rights reserved. / Page 3

The target group of this session are customers with SAP Business Warehouse systems running on SAP MaxDB, SAP partners and SAP employees who are doing BW performance analysis on SAP MaxDB systems.

As of SAP MaxDB version 7.6, MaxDB offers special features, which significantly improve the performance of SAP Business Warehouse applications.

This session provides information about these MaxDB features called **SAP MaxDB BW Feature Pack.**

The SAP Business Warehouse application itself is not part of this session.

This session is based on the internal development System S70. It is a SAP NW 7.0 system with limited BW functionality. The MaxDB Version is 7.7.07.

This session consists of a theoretical part and a second part where a performance analysis including checks of the usage of the BW Feature Pack is done directly on the system S70.

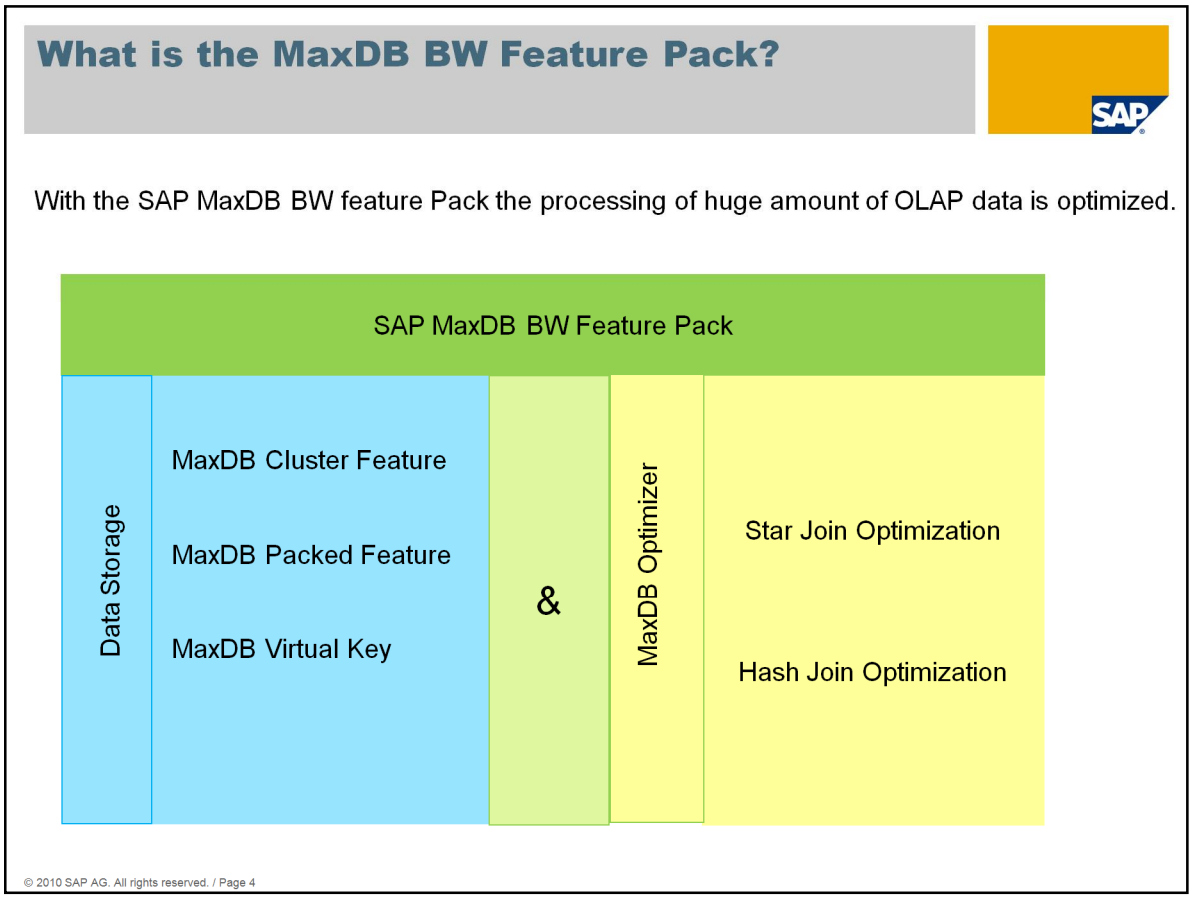

In an BW system a huge amount of data has to be processed. With the MaxDB feature pack the processing of these BW applications will be optimized.

The MaxDB BW feature pack consists of two main parts:
- the features which influence the data storage
and
- the features which affect the MaxDB optimizer.

This session gives an overview about those features, and explains how to check and how to activate them.

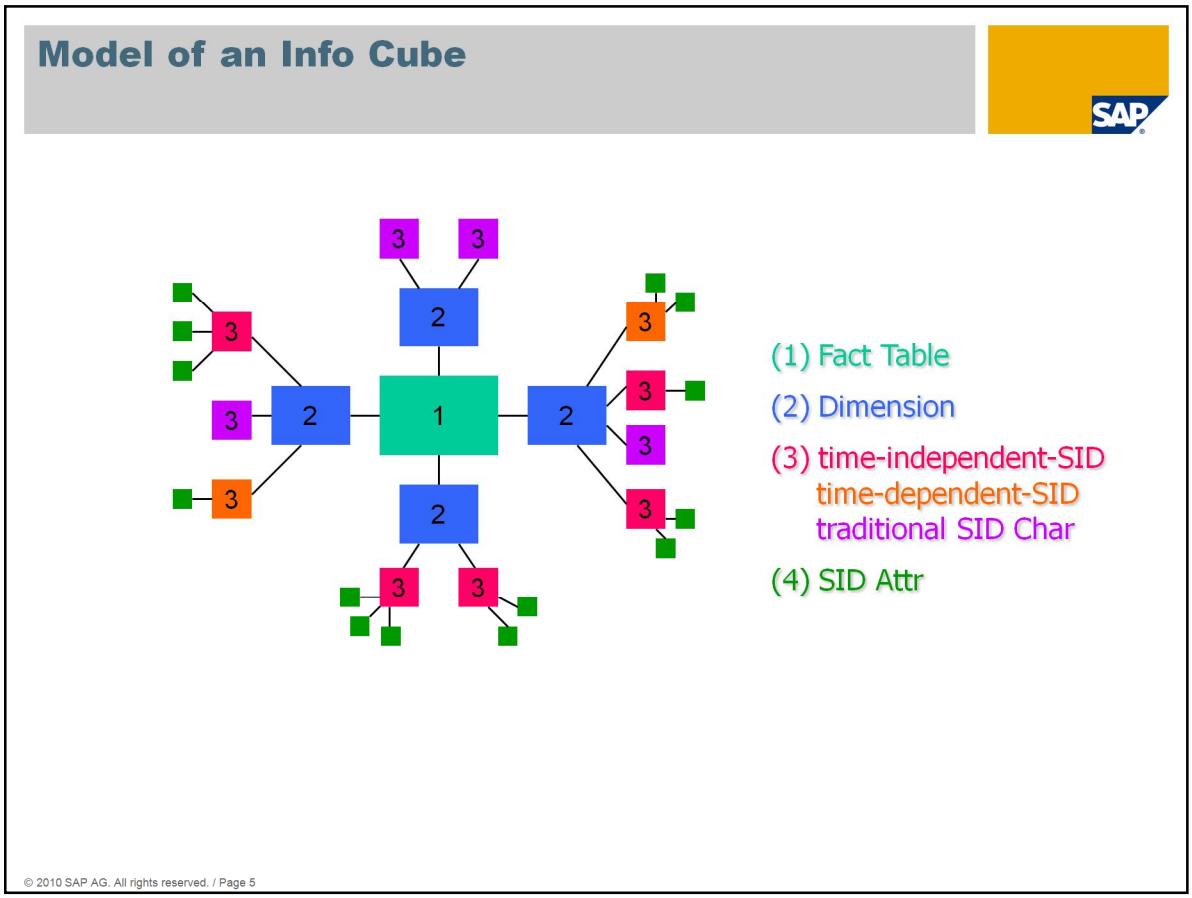

To understand the following slides a short introduction into the BW data model is necessary.

A BW system contains fact tables, dimension tables, master data tables and hierarchical tables.

BW tables created by SAP have usually the namespace /BI0. BW tables created by customers have usually the namespace /BIC. There are exceptions possible.
After the /BI0 or /BIC prefix, **fact tables** begin with '**F**' (the data is not compressed) or '**E**' (the data is compressed), e.g /BIC/**F**ZBCS_BC01.

Dimension tables begin with 'D' after the prefix, for example:
/BIC/**D**ZBCS_BC017.

After the prefix, **master data tables** begin with an '**P**' (**time-independent**) or '**Q**' (**time-dependent**) e.g. /BIC/PMaterial.
Additionally the system contains so called **SID tables** which begin with '**S**'. Those tables contain the relationship between master data values and IDs (SIDs).
The SIDs of navigation attributes are stored in tables which start with an **'X' (time-independent)** or **'Y' (time-dependent)**. for example: /BI0/**X**MATERIAL

You will find more detailed information about SAP Business Warehouse with MaxDB in **note 830468**: FAQ: Business Information Warehouse system with SAP MaxDB

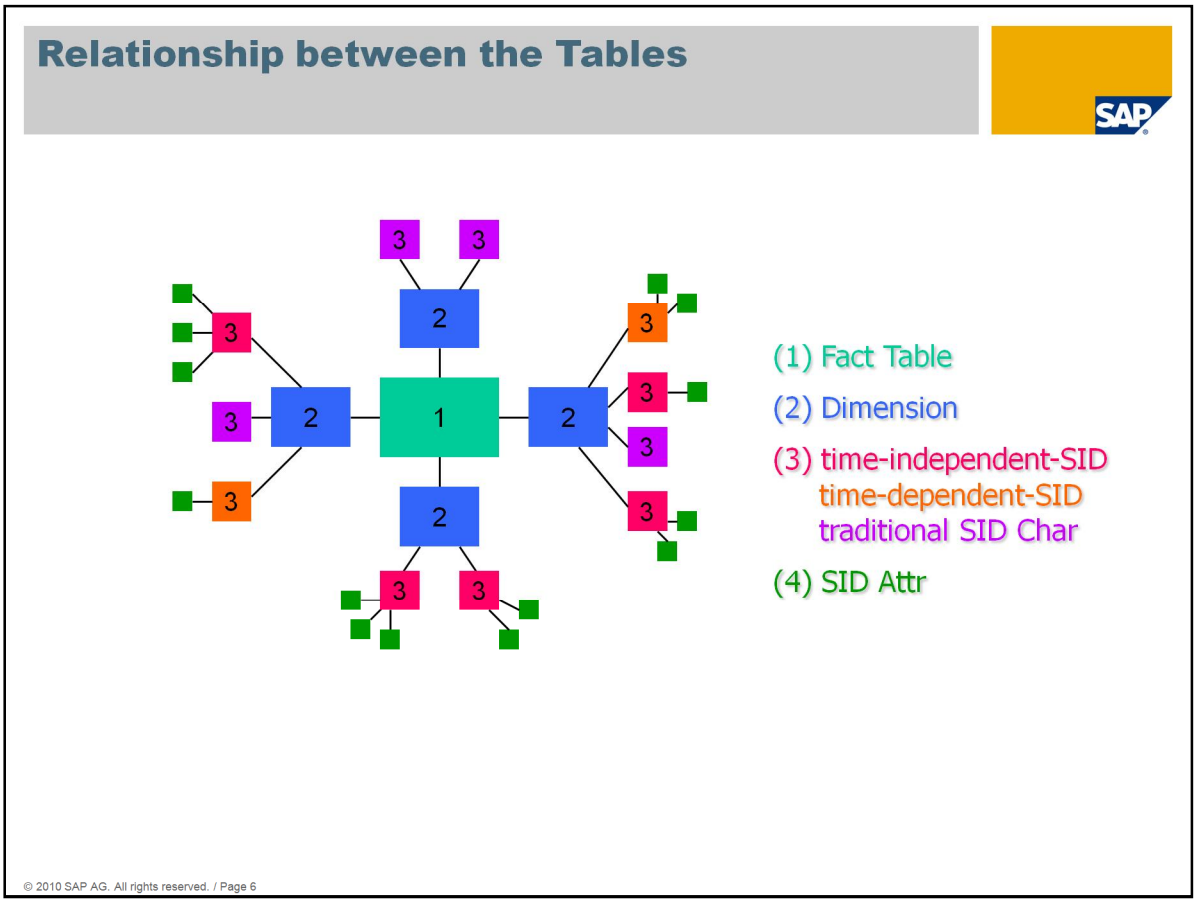

A BW query joins all data of an info cube. Mostly **one fact table** (if only one of the F or E tables is filled with data) is joined with several (**up to 16) dimension tables** and several master data tables.

If both fact tables (E and F) are filled with data, 2 SQL statements are executed. It‘s possible to have both SQL statements connected by using an UNION.

Fact Table: Contains all key fields of the associated dimension tables in its table definition. All dimension fields are indexed (KEY_...)

Dimension Tables: Column **DIMID is the primary Key** (used for the join transition to the fact table) and **all SID-columns (Master Data Ids) are indexed**.
The columns of the dimension table that form a relationship to the master data table are defined as key fields in the master data table.

Time (in)dependent Master Data (SID) Tables (prefix ‘X’)**: Column SID is the primary key** (SID is an abbreviation of the German "Stammdaten-ID").
SID tables are linked to master data tables. Attribute columns (S_...) are NOT indexed by default.

Specialty: /BIC/S… Tables

These master data tables have a primary key created on master data columns (may be several columns). There is a unique index created on the SID column.

Agenda

1. Introduction
2. MaxDB Data Storage for BW
   - Physical Table Clustering
   - Table Compression
3. Logical Table Clustering
4. MaxDB Optimizer
   - Optimized Star-Join-Processing
   - Optimized Hash Join
5. Performance Analysis and Monitoring
6. Summary

© 2010 SAP AG. All rights reserved. / Page 7

The target group of this session are customers with SAP Business Warehouse systems running on SAP MaxDB, SAP partners and SAP employees who are doing BW performance analysis on SAP MaxDB systems.

As of SAP MaxDB version 7.6, MaxDB offers special features, which significantly improve the performance of SAP Business Warehouse applications.

This session provides information about these MaxDB features called **SAP MaxDB BW Feature Pack.**

The SAP Business Warehouse application itself is not part of this session.

This session is based on the internal development System S70. It is a SAP NW 7.0 system with limited BW functionality. The MaxDB Version is 7.7.07.

This session consists of a theoretical part and a second part where a performance analysis including checks of the usage of the BW Feature Pack is done directly on the system S70.

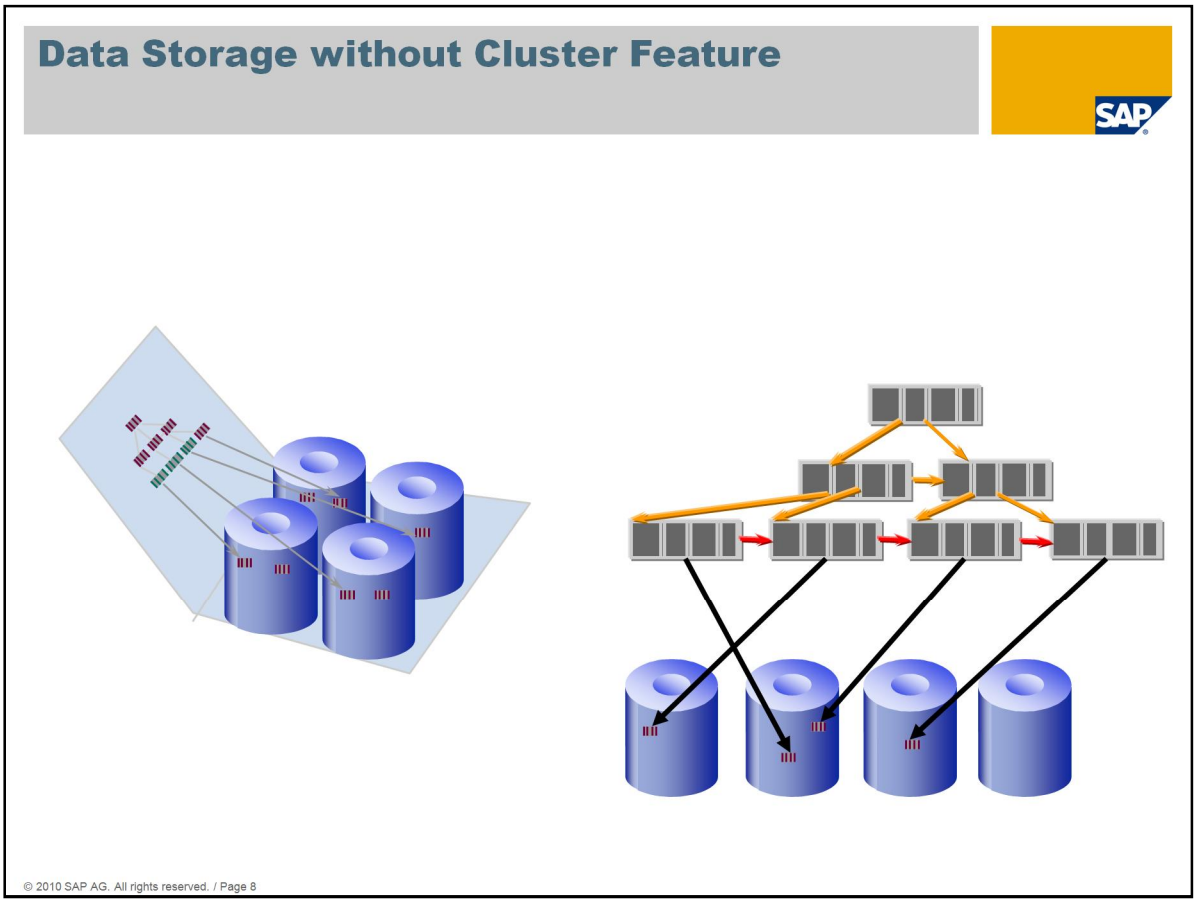

The data of a table is stored in 8 KB pages on the disk.

Pages which belong to the same table are not stored physically together but distributed over all data volumes.

Database tasks read blocks from disks in 8 KB units.

Scans don’t benefit from larger block sizes of storage systems.

To optimize the accesses to the data in a BW system other databases are using table partitioning.
This table partitioning is done in SAP BW systems via a time dimension by default.

SAP MaxDB does not offer table partitioning but the table clustering to optimize I/O accesses from disk.

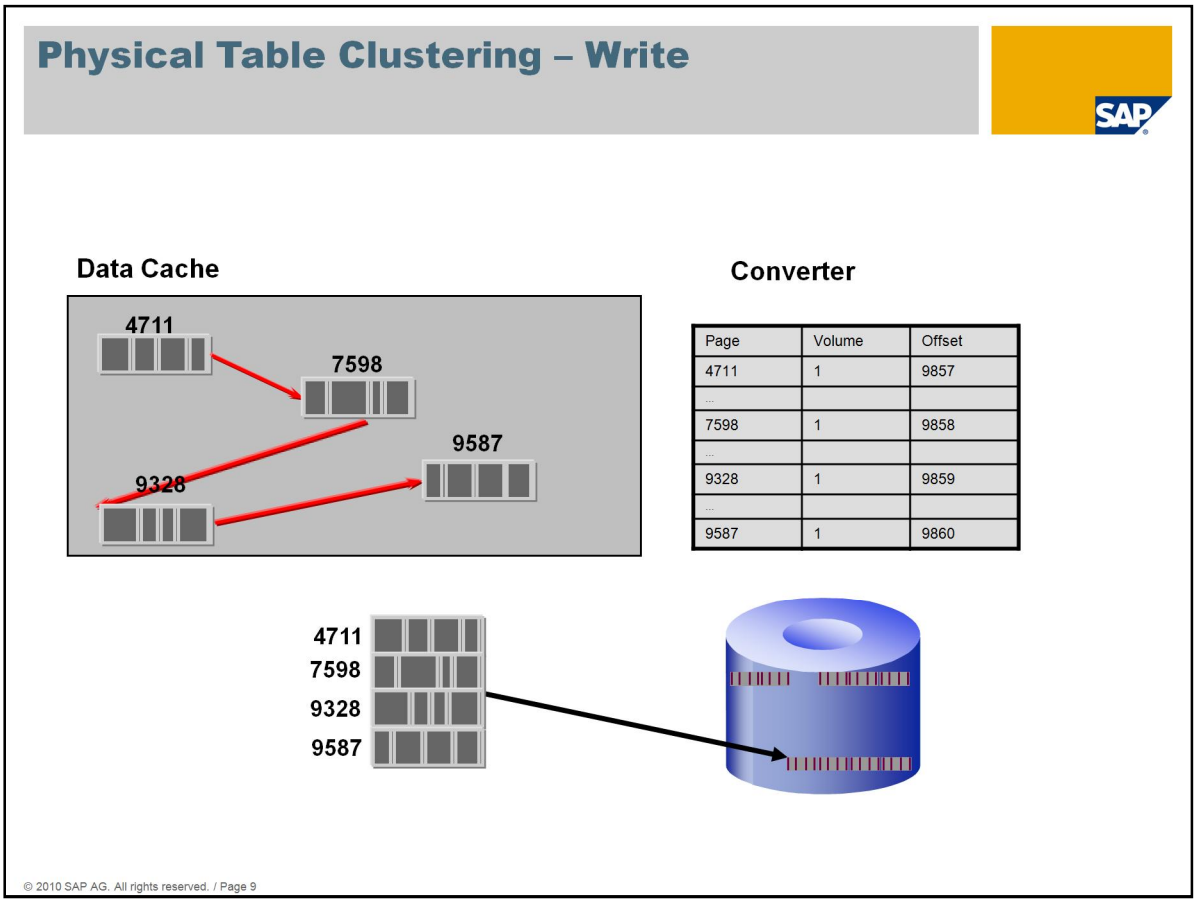

Tables can be provided with the **cluster table attribute**.

This attribute causes the system to **no longer distribute the table contents in individual blocks of 8 KB pages.**
Instead, if possible, the system stores them on the volumes **in larger blocks** in succession.

The data of clustered tables is grouped in the cache. During a savepoint, the system determines whether the data is to be stored in clusters. The data is written by pager tasks to the data volumes in segments/clusters.

The database parameter ***DataIOClusterSize*** determines the size of these segments. Before a table is clustered, you must check the setting of this parameter. The parameter should be set to at least 64 pages.

Segments are distributed across all data volumes and are stored along with data that is not clustered. On the data volumes, there are no places reserved for storing clustered data.

Physical Table Clustering – Parameter

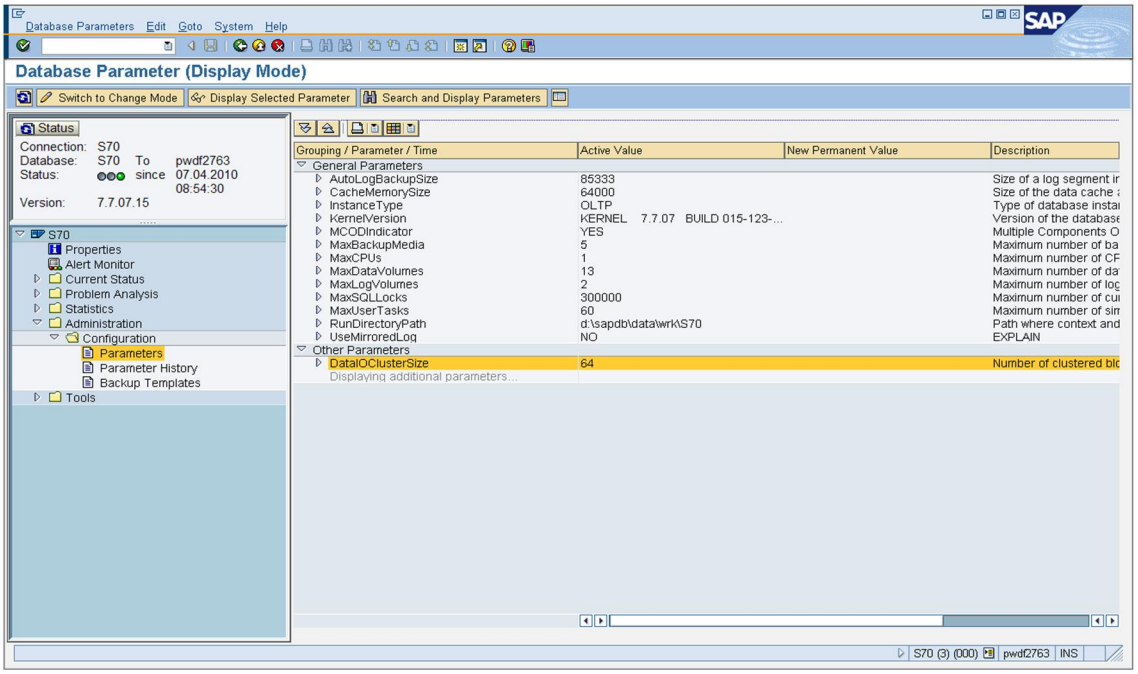

© 2010 SAP AG. All rights reserved. / Page 10

The lower and upper limits for the parameter ***DataIOClusterSize*** are:

4 <= DataIOClusterSize <= 128

The default value is 64 pages. The value depends on the I/O system you are using. For most systems a value of 64 is the best one.

Other values should be discussed with your SAP consultant/support and your hardware partner.

Physical Table Clustering – Free Blocks

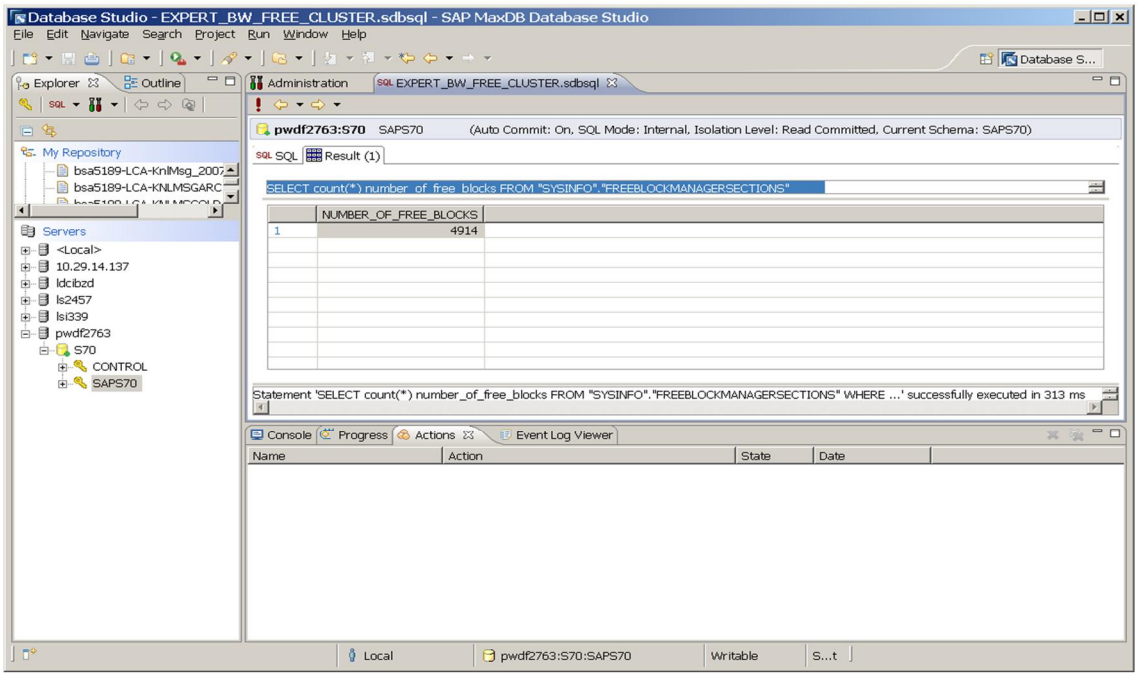

© 2010 SAP AG. All rights reserved. / Page 11

As already explained there is no reserved area on the data volumes for clustered data.

To ensure that data of a clustered table can be stored in succession, you must ensure that enough free space exists in the database to enable a good cluster grade to be obtained.

You can use the following SQL statement to determine the **number of free segments** for a segment size of 64 pages.

SAP MaxDB Version 7.6:

```
SELECT sum(sections)
FROM "SYSDD"."FREEBLOCKMANAGERSECTIONS"
WHERE freeblockcount = 64
```

SAP MaxDB Version 7.7:

```
SELECT count(*)
FROM "SYSINFO"."FREEBLOCKMANAGERSECTIONS"
WHERE freeblockcount = 64
```

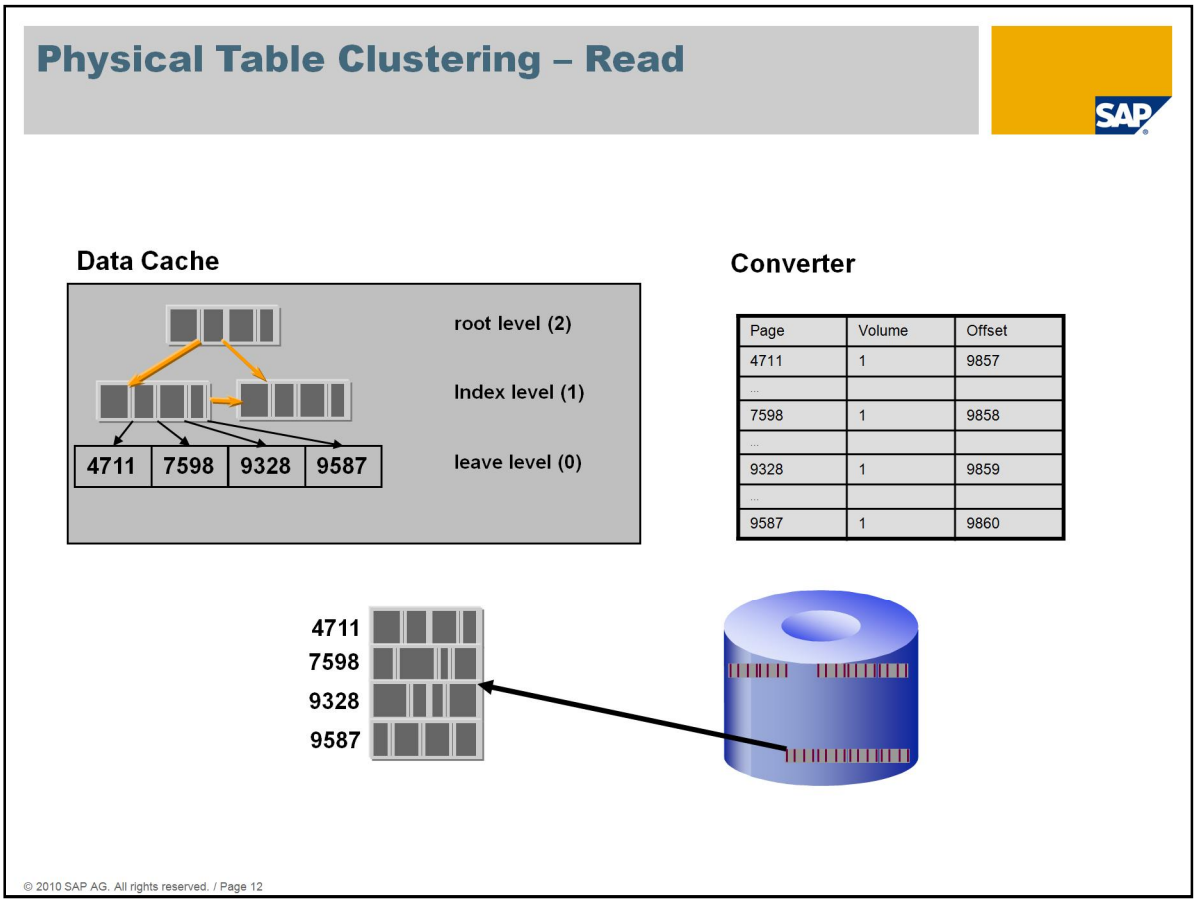

The records of a clustered table are sorted in succession on the volume and saved in segments.

During reading, they can also be read in succession. The different segments of one table are distributed across the complete data area (can be on different disks).

MaxDB reads page numbers from separators in index level (1) and checks their block positions in the converter.

The system uses one I/O operation to read blocks that are larger than 8 KB. As the data is now stored in 512 KB segments (64*8 KB), several pages can be read with one I/O call.
Experience has demonstrated that I/O times are greatly improved during sequential reading.

How many data can be read with one I/O depends on the **cluster grade (= cluster factor)** of a table. When the system reads a clustered table, it checks the cluster grade of the area that is read. If it determines that the cluster grade is poor, it marks the pages as changed and writes them in new segments to the data volumes during the next savepoint (reclustering).

You can use SELECT statements on the system tables TABLESTORAGEDETAILS and INDEXSTORAGEDETAILS to evaluate the cluster grade. For more information, see **Note *1040431 FAQ: MaxDB BW Feature Pack.***

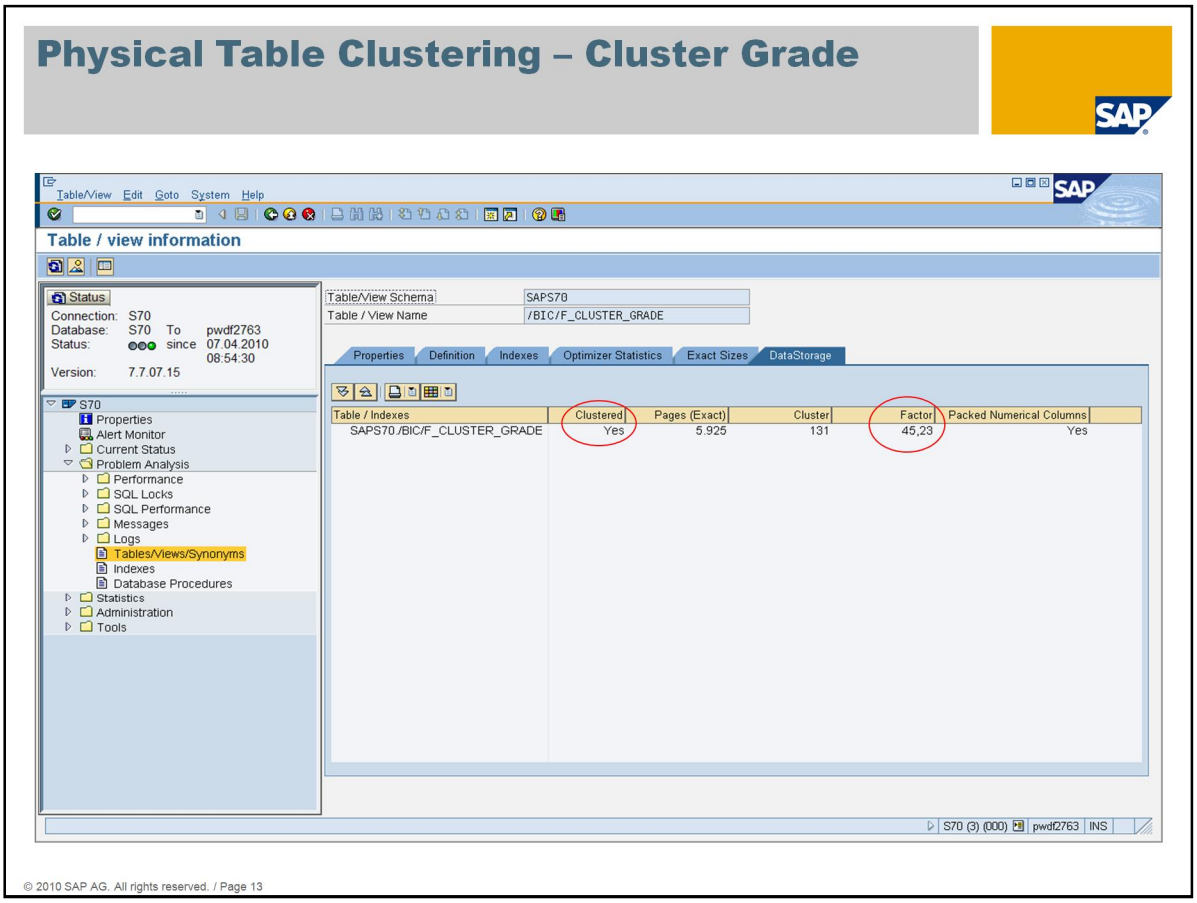

The database assistant (transaction DB50) can be used to determine the quality of the clustering.

Open *Problem analysis -> Tables/View/Synonyms* ,enter the table name and choose *Data Storage* to check the cluster attribute and cluster factor (cluster grade).

Tab *Data Storage* is only available if the table is clustered. The column *FACTOR* informs about the cluster grade. (Ratio: Pages / Number of Clusters)

The closer the cluster factor is to the value of *DataIOClusterSize*, the better is the clustering. A value of one indicates that the table is not stored in clusters.

The cluster information in the database assistant is available as of the following Basis Support Packages:

- SAP 7.10 Support Package 2
- SAP 7.00 Support Package 13
- SAP 6.40 Support Package 21
- SAP 6.20 Support Package 63
- SAP 4.6C Support Package 55

Benefit of Physical Table Clustering

- Database tasks read blocks from disks in cluster units
- Scans benefit from larger block sizes of storage systems
- Field experience: Scans are 5-6 times faster in average

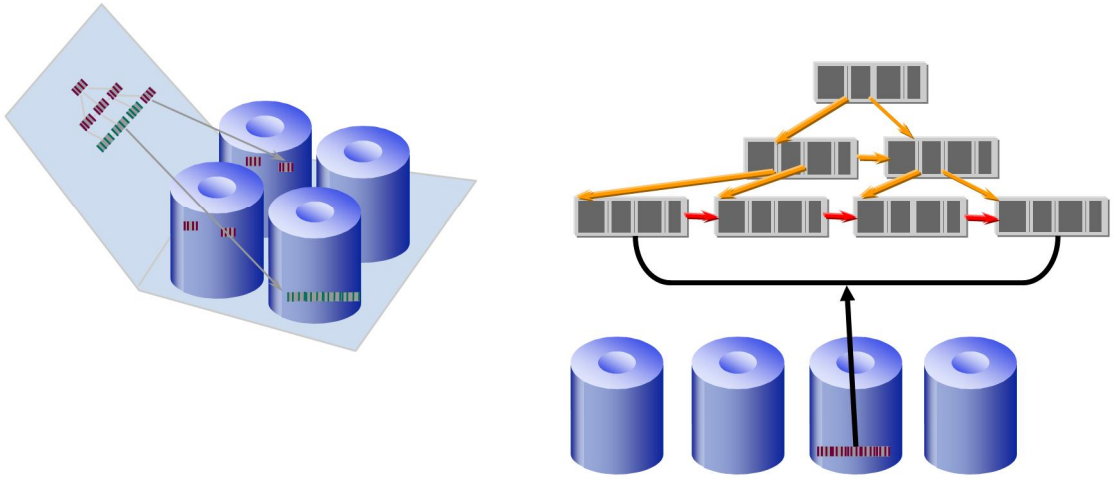

© 2010 SAP AG. All rights reserved. / Page 14

More information about the cluster feature can be found in **note:**

1264776 FAQ: SAP MaxDB cluster table attribute

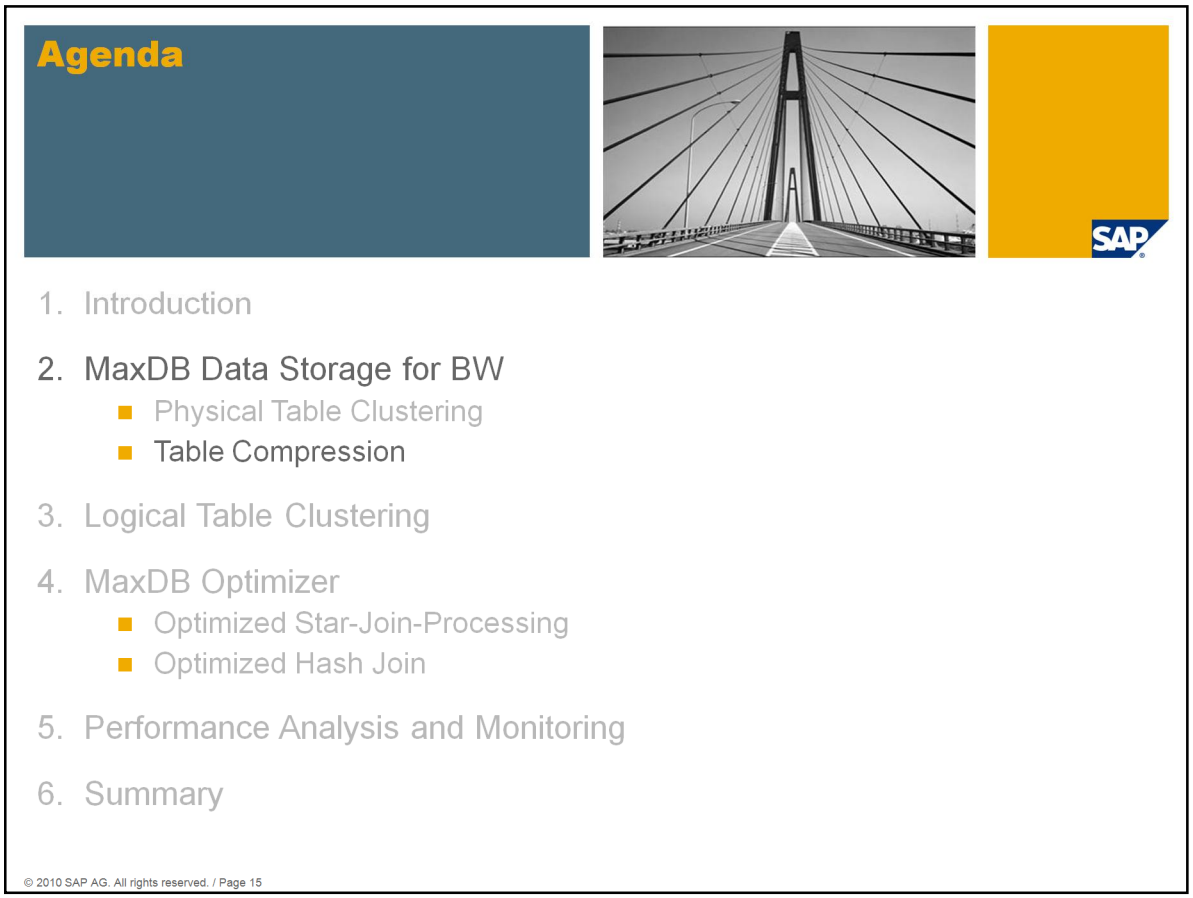

The target group of this session are customers with SAP Business Warehouse systems running on SAP MaxDB, SAP partners and SAP employees who are doing BW performance analysis on SAP MaxDB systems.

As of SAP MaxDB version 7.6, MaxDB offers special features, which significantly improve the performance of SAP Business Warehouse applications.

This session provides information about these MaxDB features called **SAP MaxDB BW Feature Pack.**

The SAP Business Warehouse application itself is not part of this session.

This session is based on the internal development System S70. It is a SAP NW 7.0 system with limited BW functionality. The MaxDB Version is 7.7.07.

This session consists of a theoretical part and a second part where a performance analysis including checks of the usage of the BW Feature Pack is done directly on the system S70.

Table Compression – Packed

- Compression of numeric columns only
  - In general fact tables have only numeric columns with a maximum value length of 15 digits (FIXED)
  - Two digits can be stored in 7 Bit instead of 8 Bit
  - Length indicator can be stored in 4 Bit instead of 8 Bit

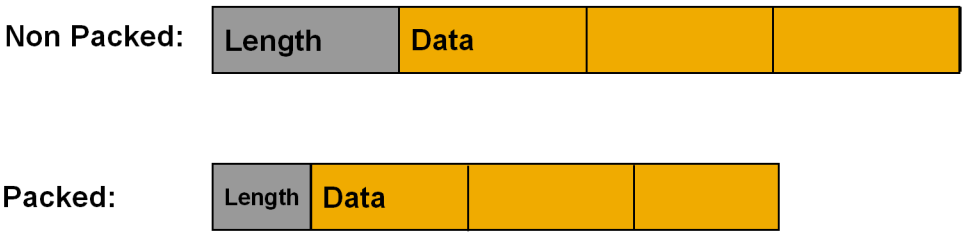

© 2010 SAP AG. All rights reserved. / Page 16

The PACKED table attribute causes integer values to be compressed so that they require less space on the volumes. This results in a faster I/O access and reduced displacement in the data cache.

You can check whether a table has the PACKED attribute by looking at the PACKEDNUMBERCOLUMNS column in the system table FILES or using transaction DB50.

Primary key columns are never packed.

Precondition that columns can be packed:

•All non primary key columns are numeric.

•The columns are defined as NOT NULL.

•Each column contains values with a maximum length of 15 digits.

The PACKED attribute will be ignored if the preconditions are not fulfilled.

Benefit: Field experience: 60-70% smaller tables

Later in this session there is a comparison of the table size before and after a table has been packed.

Table Compression – Packed

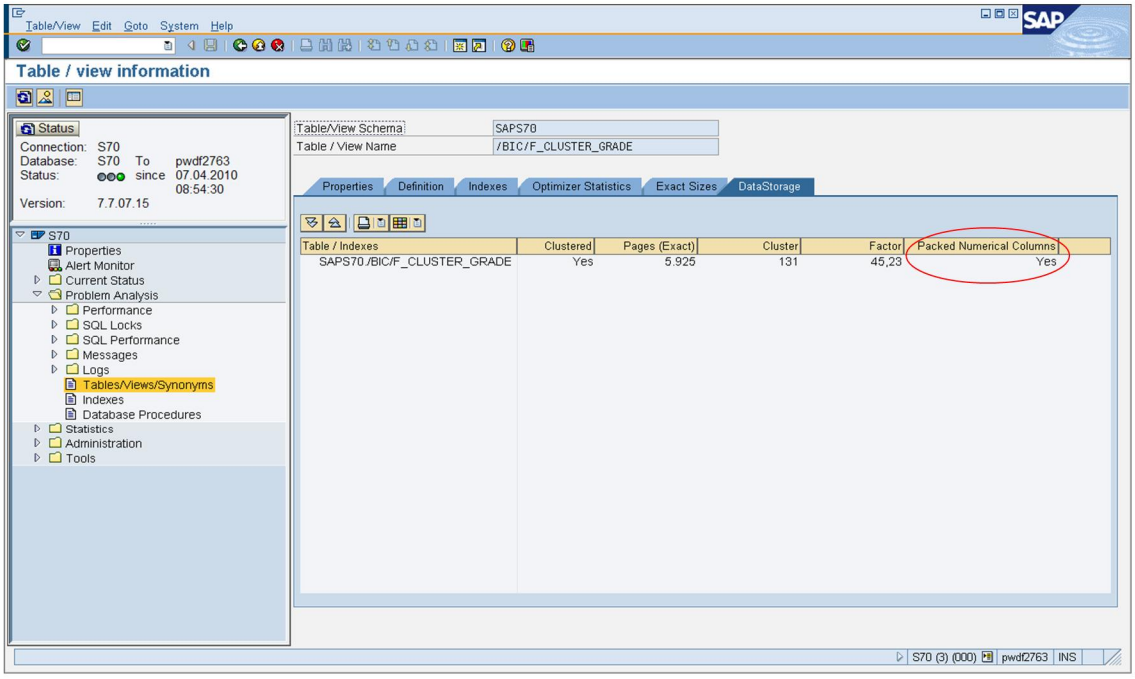

© 2010 SAP AG. All rights reserved. / Page 17

The Database Assistant (transaction DB50) can be used to determine if a fact table is packed.

Open *Problem analysis -> Tables/View/Synonyms,* enter the table name and chose *Data Storage* to check the output of *Packed Numeric Columns .*

Precondition that the Tab *Data Storage* is listed is that the table has the cluster attribute.

Agenda

1. Introduction
2. MaxDB Data Storage for BW
   - Physical Table Clustering
   - Table Compression
3. Logical Table Clustering
4. MaxDB Optimizer
   - Optimized Star-Join-Processing
   - Optimized Hash Join
5. Performance Analysis and Monitoring
6. Summary

© 2010 SAP AG. All rights reserved. / Page 18

The target group of this session are customers with SAP Business Warehouse systems running on SAP MaxDB, SAP partners and SAP employees who are doing BW performance analysis on SAP MaxDB systems.

As of SAP MaxDB version 7.6, MaxDB offers special features, which significantly improve the performance of SAP Business Warehouse applications.

This session provides information about these MaxDB features called **SAP MaxDB BW Feature Pack.**

The SAP Business Warehouse application itself is not part of this session.

This session is based on the internal development System S70. It is a SAP NW 7.0 system with limited BW functionality. The MaxDB Version is 7.7.07.

This session consists of a theoretical part and a second part where a performance analysis including checks of the usage of the BW Feature Pack is done directly on the system S70.

Virtual Key Concept
Create Table Command

- Use time characteristic internally as part of virtual key

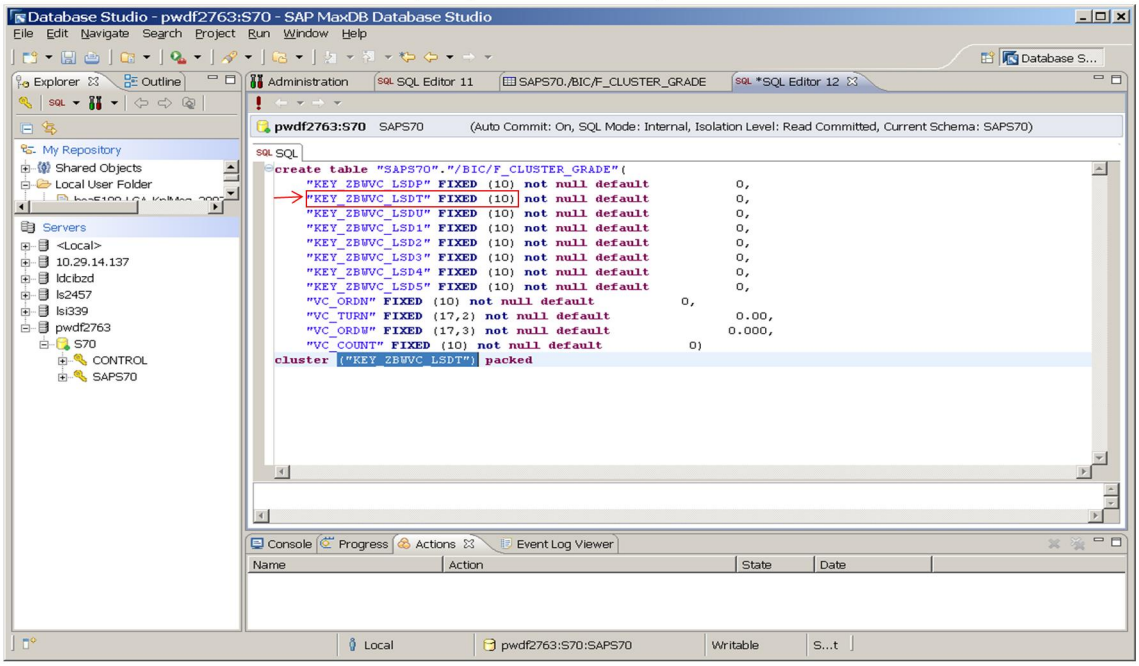

© 2010 SAP AG. All rights reserved. / Page 19

Fact tables do not have a user defined primary key. To optimize accesses on huge tables like fact tables a **virtual key concept** (logical table clustering) has been implemented.

A virtual key allows you to create a unique key for **non-unique columns**. From a technical point of view, this virtual key consists of a non-unique key field (KEY_ZBWVC_LSDT) and a unique sequential number (SEQUENCE).

If a table gets a virtual key, the content of the table is sorted by the virtual key (logical table clustering). Additionally the data is also stored in clusters (physical table clustering) automatically.

Usually the data will be promptly loaded from the source system into the BW system. Therefore new data will have different date values than the already existing data. This is in MaxDB systems the reason for using the **time characteristic** for the virtual key definition.

The virtual key as well as the physical clustering of the **fact tables** are created in a BW system during the table creation.

Table Definition

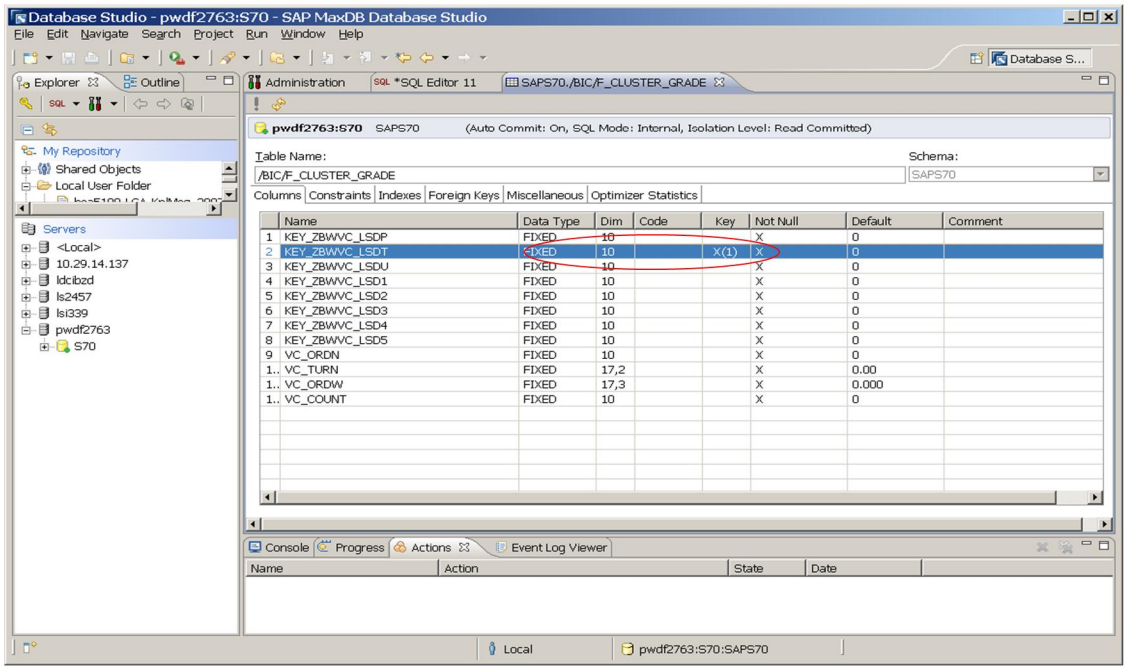

© 2010 SAP AG. All rights reserved. / Page 20

In BW systems the virtual key concept is used for fact tables which don't have a user defined primary.

By default the time characteristic is used for virtual cluster key.
The table definition (system table DOMAIN.COLUMNS) shows the time dimension as the virtual key.

The sequence which is implicitly added to each time dimension entry is not visible.

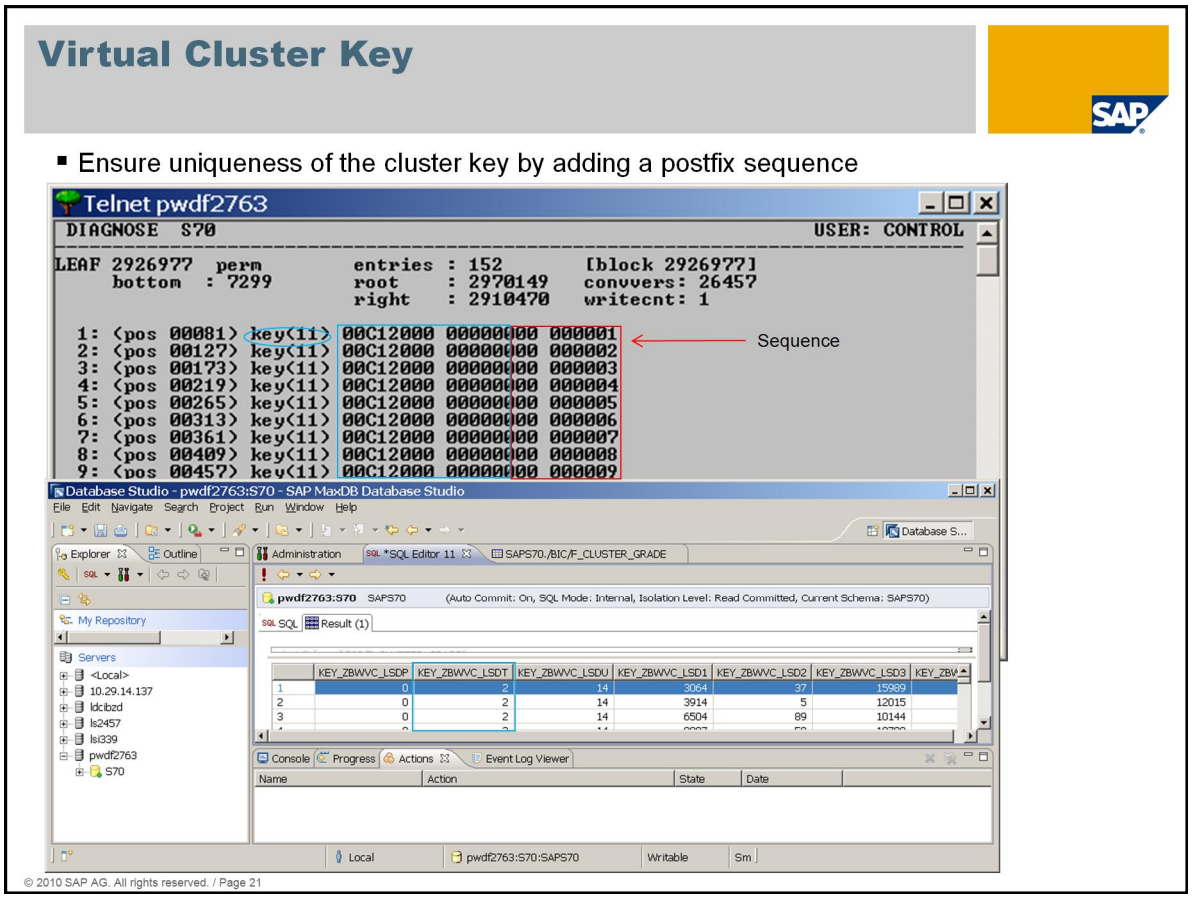

Benefit:

•**Join via virtual cluster key** instead of join via index

•Field experience: **50% faster join** (precondition: without IO)

During the installation of BW 7.0 with MaxDB 7.7 the feature package is automatically activated. New cube and aggregate tables (fact tables) will be created automatically with BW Feature Pack attributes (clustered, packed).

During an upgrade fact tables which are NOT clustered won't be converted implicitly. These fact tables have to be converted with the report RSDU_CLUSTER_FACT_ADA.

Later in this session you will get the information how to check these table attributes and how to cluster already existing tables in the system.

The sequence number of the virtual key is invisible for the application.
It can be seen only on page level with a database diagnosis tool.

Agenda

1. Introduction
2. MaxDB Data Storage for BW
   - Physical Table Clustering
   - Table Compression
3. Logical Table Clustering
4. MaxDB Optimizer
   - Optimized Star-Join-Processing
   - Optimized Hash Join
5. Performance Analysis and Monitoring
6. Summary

© 2010 SAP AG. All rights reserved. / Page 22

The target group of this session are customers with SAP Business Warehouse systems running on SAP MaxDB, SAP partners and SAP employees who are doing BW performance analysis on SAP MaxDB systems.

As of SAP MaxDB version 7.6, MaxDB offers special features, which significantly improve the performance of SAP Business Warehouse applications.

This session provides information about these MaxDB features called **SAP MaxDB BW Feature Pack.**

The SAP Business Warehouse application itself is not part of this session.

This session is based on the internal development System S70. It is a SAP NW 7.0 system with limited BW functionality. The MaxDB Version is 7.7.07.

This session consists of a theoretical part and a second part where a performance analysis including checks of the usage of the BW Feature Pack is done directly on the system S70.

Optimized Star Join

- The MaxDB optimizer is able to identify a fact table
- Join from the first dimension cloud to the fact table

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

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

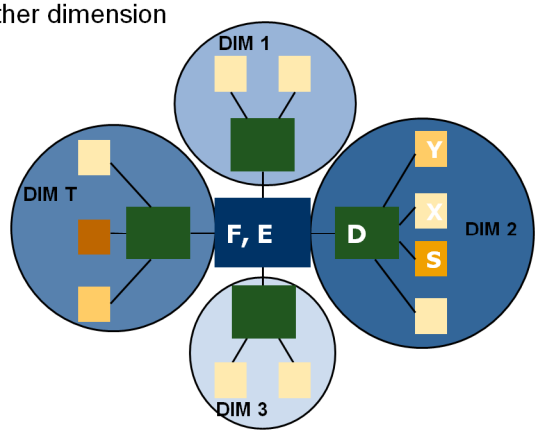

- **Benefit:**
  Prevent long runtimes of joins due to undesirable joins between dimension clouds

© 2010 SAP AG. All rights reserved. / Page 23

The Star Join Optimization of MaxDB is a special type of search for the best join sequence.

MaxDB does **not use** any special execution strategies such as **bitmap indexes**.

To find the best access strategy, the optimizer divides the query into individual joins. An individual join consists of a dimension table and all tables that are directly linked with this dimension cloud.

The best execution plan is then determined for each of these individual joins (dimension cloud). The system then searches for the best sequence for all dimension clouds and the fact table. Using this method and knowledge of the BI data model (Note 830468), the search area is kept as small as possible.

Optimized Star Join
Parameter Configuration

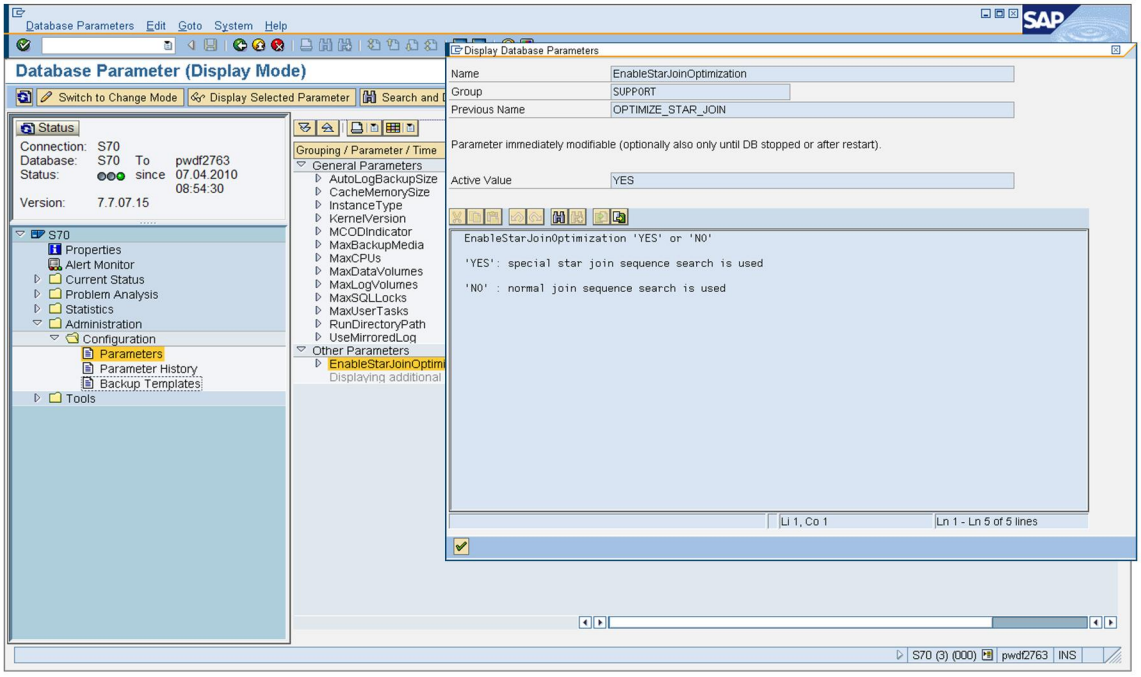

© 2010 SAP AG. All rights reserved. / Page 24

Prerequisits:

- MaxDB Version 7.6.01 Build 7 or higher.
- MaxDB Version 7.7.07 or higher.
- Check/Set the parameter *EnableStarJoinOptimization* to the value YES

Optimized Star Join
Attribute *Fact*

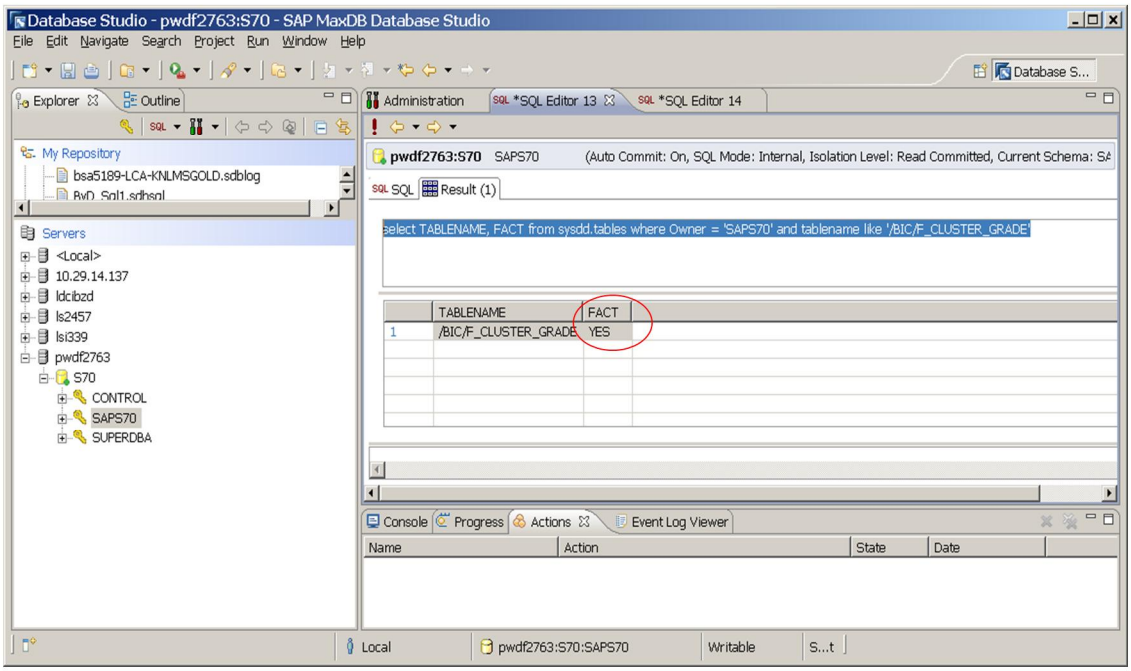

© 2010 SAP AG. All rights reserved. / Page 25

The MaxDB Optimizer can execute the Star Join Optimization only if fact tables are marked as fact table in the database catalog.

During performance analysis the so called *fact* attribute should be checked.

Unfortunately it is not possible to check the fact attribute via CCMS.

The fact attribute can be checked via select on sysdd.tables in Database Studio only.

```
select tablename, fact
from sysdd.tables
where tablename = '<tablename>'
```

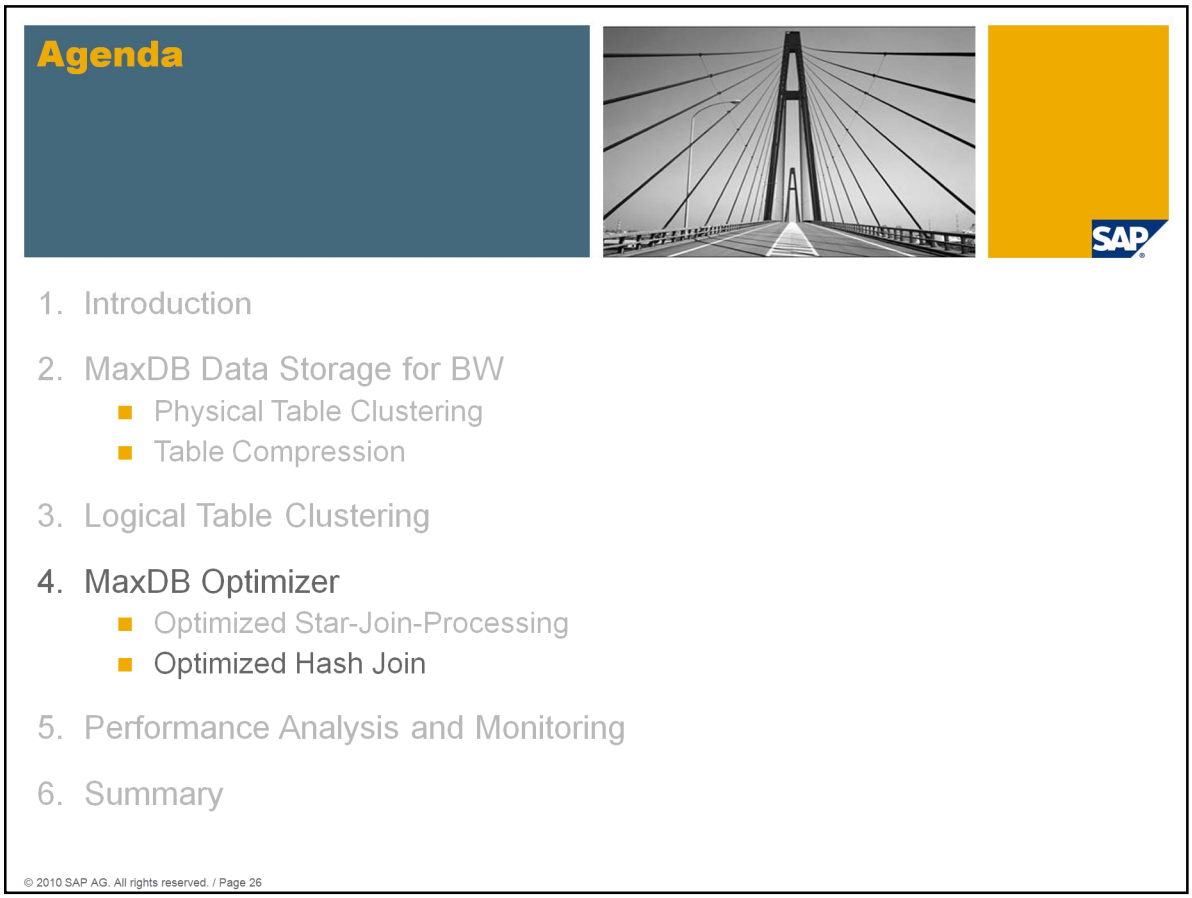

The target group of this session are customers with SAP Business Warehouse systems running on SAP MaxDB, SAP partners and SAP employees who are doing BW performance analysis on SAP MaxDB systems.

As of SAP MaxDB version 7.6, MaxDB offers special features, which significantly improve the performance of SAP Business Warehouse applications.

This session provides information about these MaxDB features called **SAP MaxDB BW Feature Pack.**

The SAP Business Warehouse application itself is not part of this session.

This session is based on the internal development System S70. It is a SAP NW 7.0 system with limited BW functionality. The MaxDB Version is 7.7.07.

This session consists of a theoretical part and a second part where a performance analysis including checks of the usage of the BW Feature Pack is done directly on the system S70.

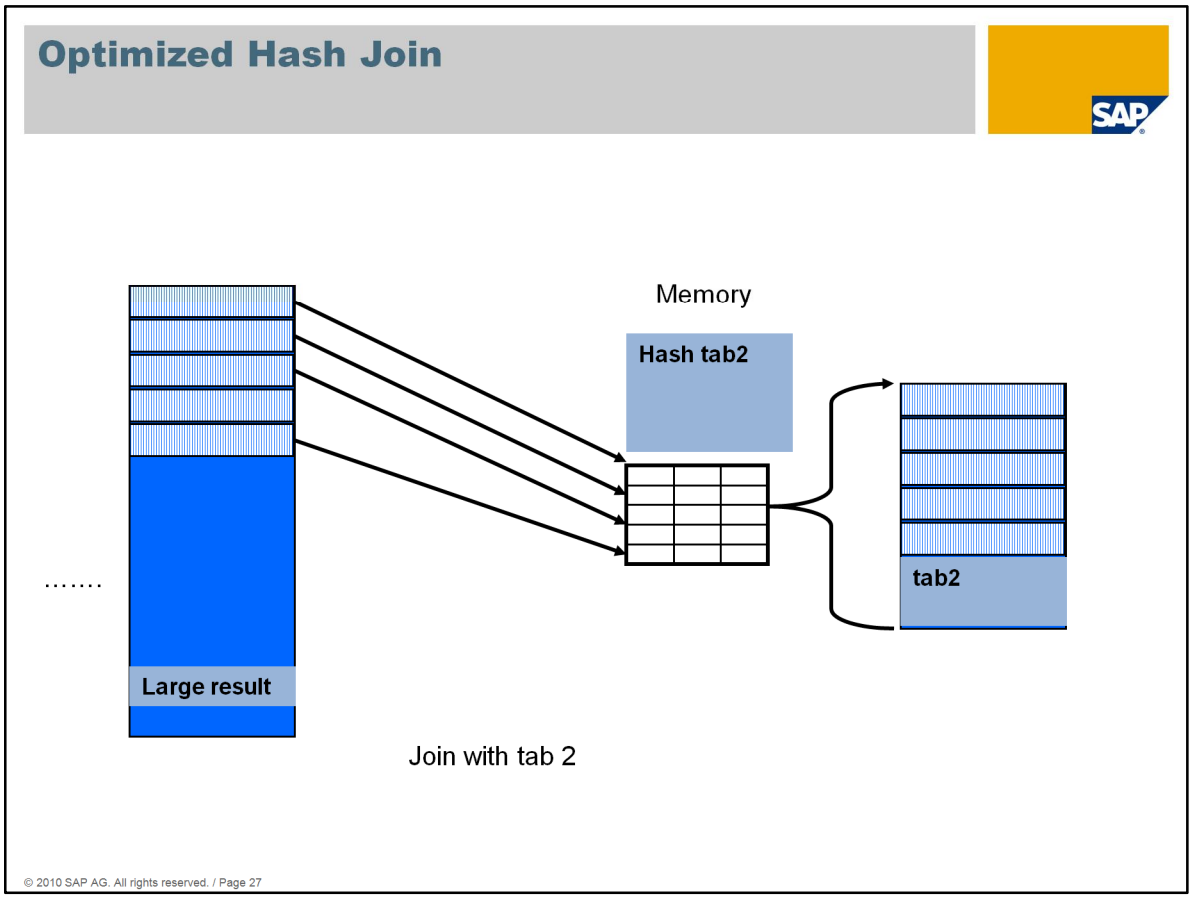

This slide shows an example of a hash join processing. When the hash join is used, the temporary result of a join step is created as hash table in the memory. When the hash join is not used the temporary result is created in the data cache as a b*-tree.

Joins via hash tables make sense, if the ratio between the size of an intermediate join result and the size of the next join table exceeds a certain value.
This Hash Join Optimization is switched on by default and can be deactivated with the parameter *EnableJoinHashTableOptimization*.

In the explain output of an SQL statement you will find the information *table hashed* when the hash join is used.

Benefit:
For each 1 Mio rows in a result set SQL statements are 4-5 seconds faster.
Field experience shows, that selects are 2 times faster.

Optimized Hash Join
Parameter Configuration

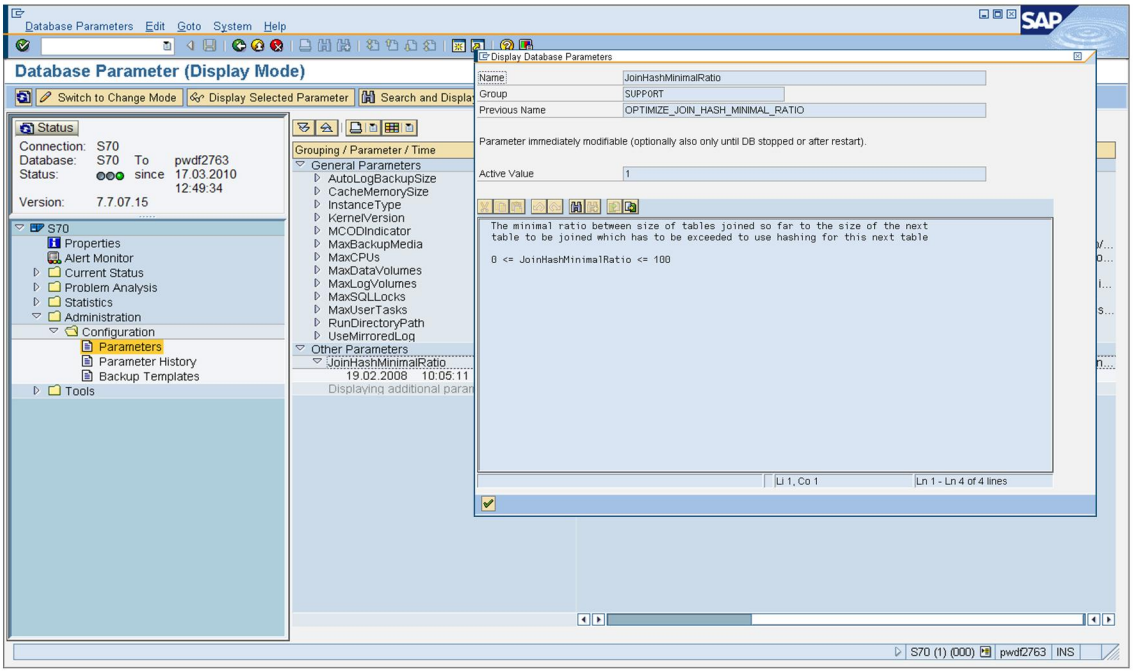

© 2010 SAP AG. All rights reserved. / Page 28

The parameter *JoinHashMinimalRatio* (called OPTIMIZE_JOIN_HASH_MINIMAL_RATIO in MaxDB version 7.6) defines the minimal ratio for the usage of hash join.

The default value of this parameter is 1%.

Agenda

1. Introduction
2. MaxDB Data Storage for BW
   - Physical Table Clustering
   - Table Compression
3. Logical Table Clustering
4. MaxDB Optimizer
   - Optimized Star-Join-Processing
   - Optimized Hash Join
5. Performance Analysis and Monitoring
6. Summary

© 2010 SAP AG. All rights reserved. / Page 29

The target group of this session are customers with SAP Business Warehouse systems running on SAP MaxDB, SAP partners and SAP employees who are doing BW performance analysis on SAP MaxDB systems.

As of SAP MaxDB version 7.6, MaxDB offers special features, which significantly improve the performance of SAP Business Warehouse applications.

This session provides information about these MaxDB features called **SAP MaxDB BW Feature Pack.**

The SAP Business Warehouse application itself is not part of this session.

This session is based on the internal development System S70. It is a SAP NW 7.0 system with limited BW functionality. The MaxDB Version is 7.7.07.

This session consists of a theoretical part and a second part where a performance analysis including checks of the usage of the BW Feature Pack is done directly on the system S70.

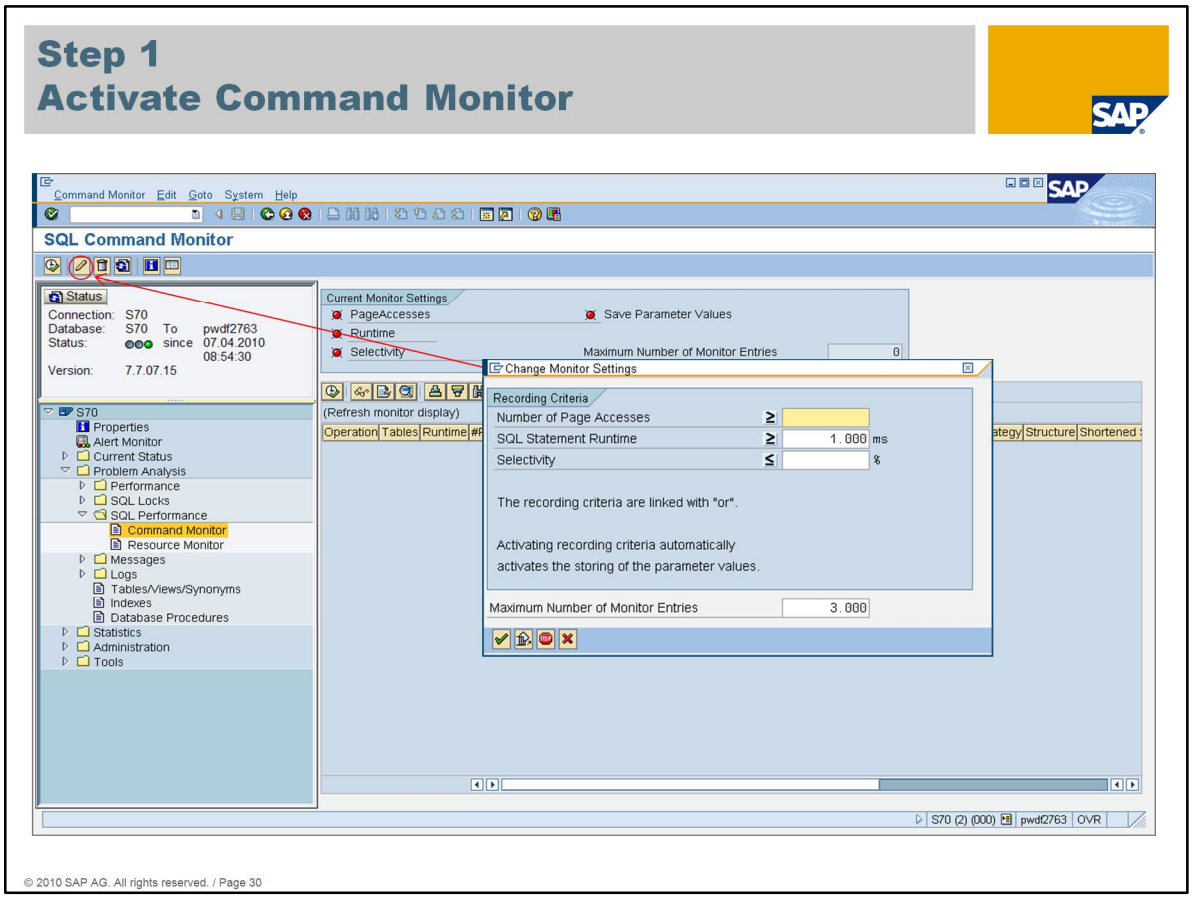

The MaxDB *command monitor* is used to find those queries which have a long runtime and should be optimized.

Before starting the performance analysis of BW Queries the *command monitor* has to be switched on.

For this session we use the default values.

When there is high workload on a productive system the default values should be changed to avoid that too many statements are logged in the *command monitor*. The new values depend on the response time of the queries which should be analyzed.

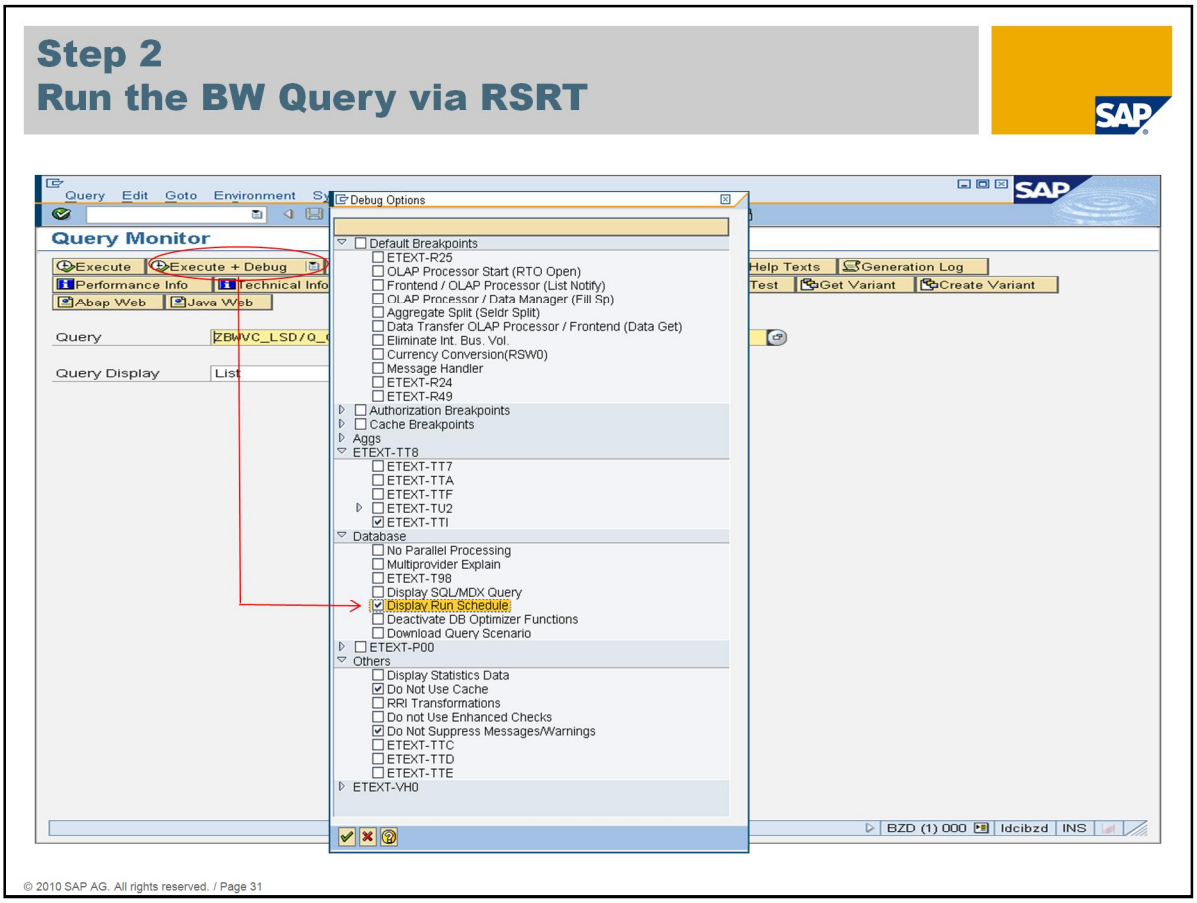

In a SAP Business Warehouse system transaction RSRT can be used to run the BW-Queries which are causing performance problems.

Choose the *Execute + Debug* button and mark the option *Display Run Schedule.* Then the explain plan after the query has been executed is displayed.

BW-Traces can also be created with transaction RSRTRACE. Advantage of these traces are that the SAP support can use these traces to reproduce the performance problem with debug (*Display Run Schedule*) option.

For this session in System S70 a Report *Z_EXPERT_BW_1* is directly started from transaction SE38.

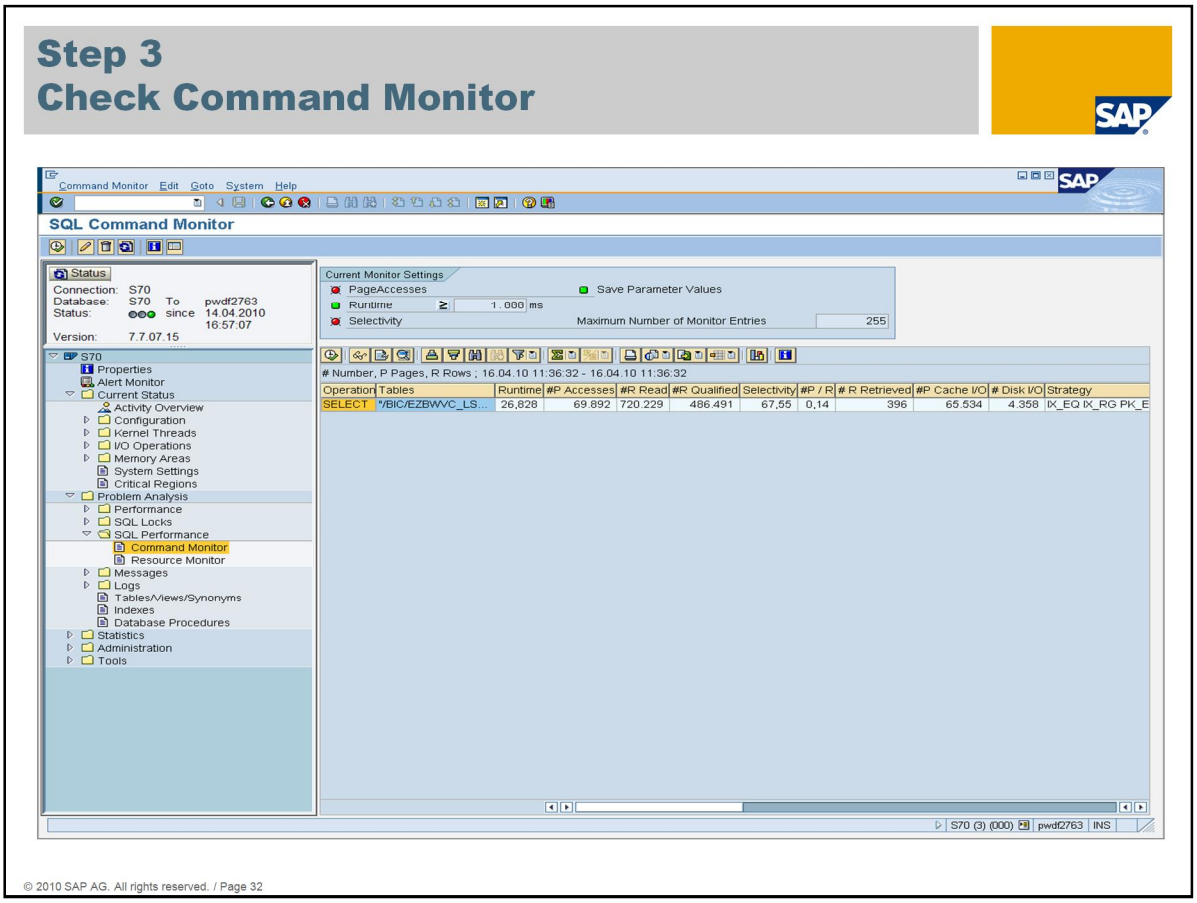

After the problematic query has been executed, the command monitor can be checked for long running statements.
Those commands which have a long runtime are relevant for a detailed performance analysis.

Reasons for a long runtime in BW systems can be:

- A huge amount of data that has to be read.
- Lots of I/O has to be done to get the data.
- Not optimized accesses to the data.
- Not the best possible join transition is used.
- The hash join optimization is not used.

In this example the statement ran more than 20 seconds and nearly 70.000 page accesses with more than 4.000 disk I/O had to be done.

It has to be checked why so many page and disk accesses were done. Which strategy is used to access the data?

With a double click on the SQL statement a detailed analysis is possible.

Step 4
Detailed SQL Statement analysis

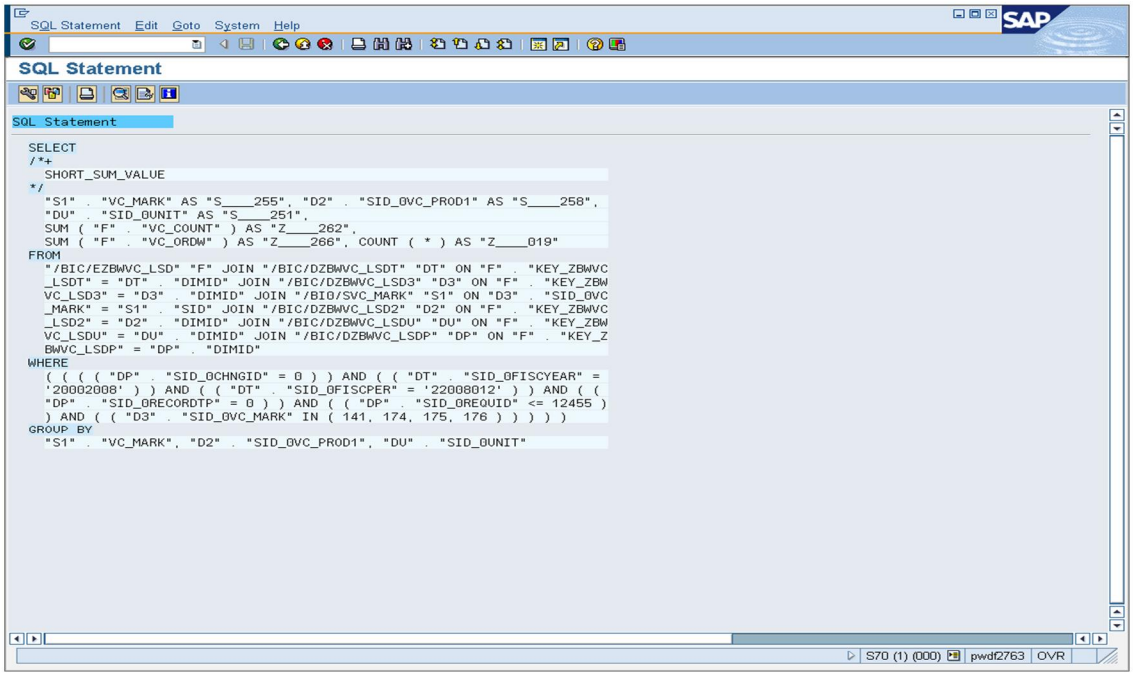

© 2010 SAP AG. All rights reserved. / Page 33

The first part of the SQL command analysis should be the analysis of the SQL statement itself.

You must know:

- what's the name of the fact table
- which dimension tables are involved in this SQL statement
- are there any master data tables involved in this SQL statement

To get a better overview of the SQL statement SAP recommends to use the *Database Studio* to do this SQL statement structure analysis.

Copy the SQL statement above and paste it into the *Database Studio SQL Editor* to get the SQL statement formatted.

Step 5
SQL Command structure analysis

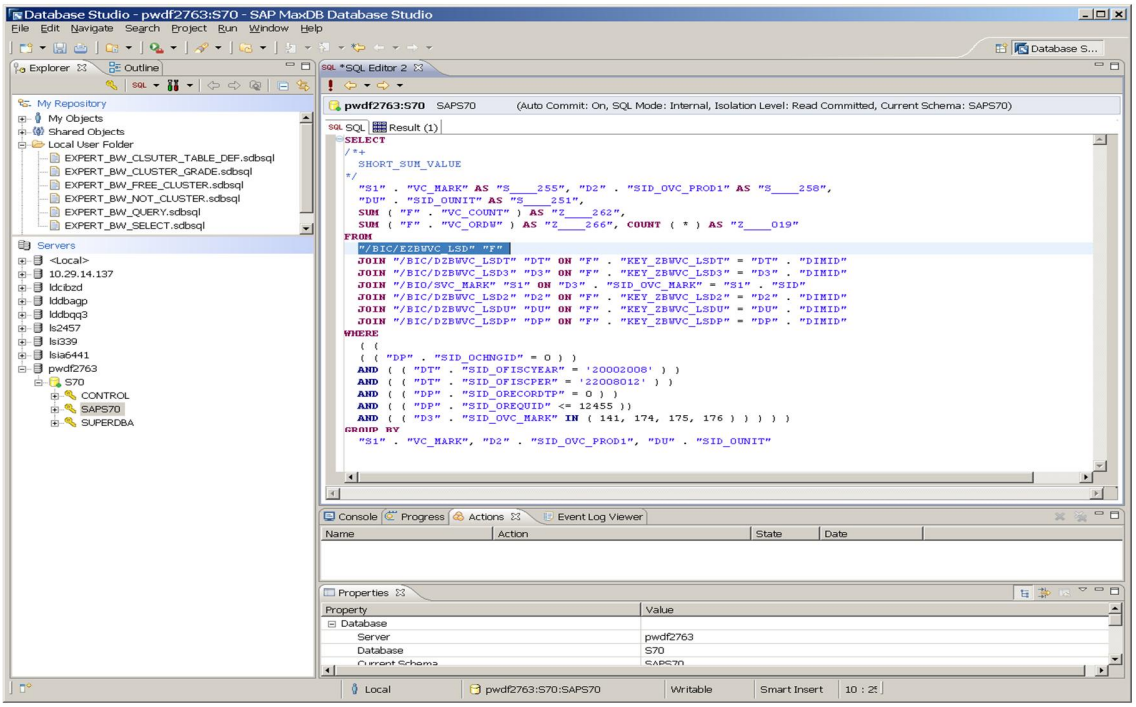

© 2010 SAP AG. All rights reserved. / Page 34

The SQL structure analysis of a BW-Query starts with finding out the name of the fact table. In this example the name of the fact table is */BIC/**E**ZBWVC_LSD* (Remember the name convention: E = compressed fact table).

You must know this name if you want to find out to which Cube the fact table belongs to.

The join transition (FROM part) shows the relation between the involved dimension and master data tables to the fact table.

In this example 5 dimension tables are joined with the fact table. */BIC/DZBWVC_LSDT* (Time dimension); */BIC/DZBWVC_LSDP* (package dimension); */BIC/DZBWVC_LSDU* (unit dimension); */BIC/DZBWVC_LSD2* and */BIC/DZBWVC_LSD3* (user defined dimensions)

Addionally dimension */BIC/DZBWVC_LSD3* is joined with the master data SID table */BI0/SVC_MARK*.

Keep this information in mind, when you have a closer look to the explain output in the next step.

Step 6
SQL Command analysis - EXPLAIN

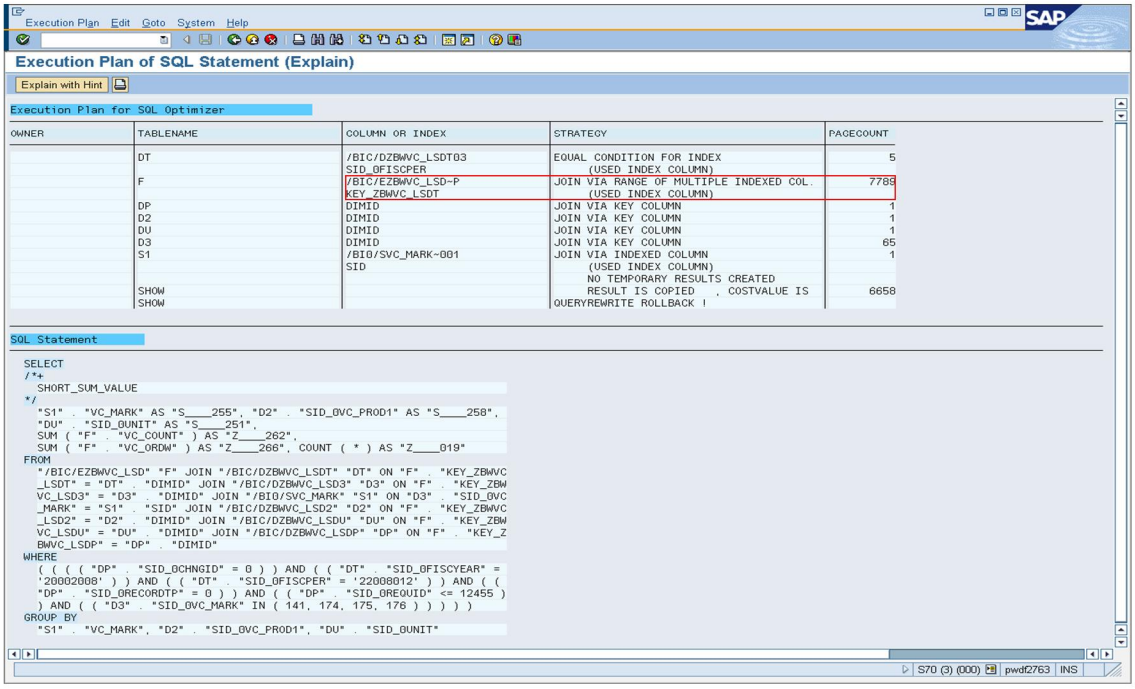

© 2010 SAP AG. All rights reserved. / Page 35

To decide if the SQL command is executed in the best way a closer look to the local predicates (where condition) has to be done, too.

This SQL command has the following local predicates:
DT (alias of /BIC/DZBWVC_LSDT): SID_0FISCPER" = '22008012' and SID_0FISCYEAR = '20002008'

DP (alias of /BIC/DZBWVC_LSDP): SID_0CHNGID = 0 and SID_0RECORDTP = 0 and SID_0REQUID <= 12455

D3 (alias of /BIC/DZBWVC_LSD3): SID_0VC_MARK IN (141, 174, 175, 176)

The optimizer chooses DT as the first table, which looks good. The second table is the fact table, which sounds good as well. Then the other tables follow.

Note: When the feature pack is used, the system starts in most cases with the time dimension and the join transition is processed by using the virtual key (strategy: primary key access). In this example the strategy JOIN VIA RANGE OF MULTIPLE INDEXED COL. via index /BIC/EZBWVC_LSD~P is used. In this case the physical clustering of the fact table data has no affect while the data is read.

The following slide explains why the cluster feature cannot be used.

JOIN VIA RANGE OF MULTIPLE INDEXED COL

Select time dimension via Equal condition for index

Join fact table via range of multiple index columns

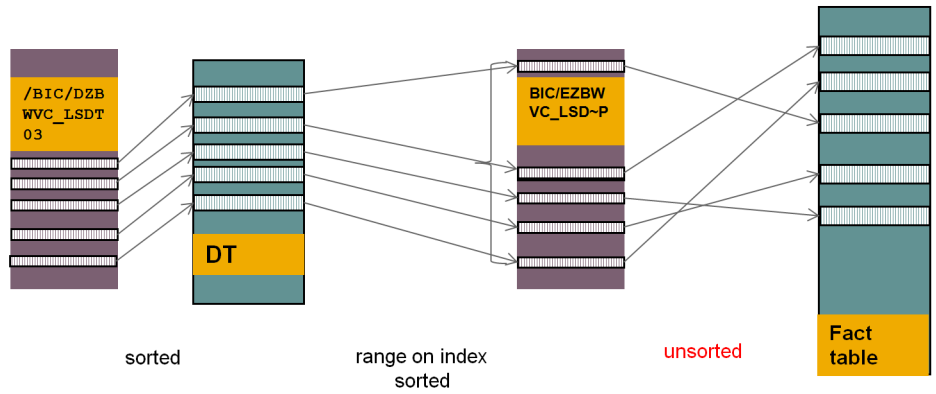

No benefit of the cluster feature !

© 2010 SAP AG. All rights reserved. / Page 36

This slide shows the first execution steps of the BW-Query in this session.

Starting with the equal access on index /BIC/DZBWVC_LSDT03 of the time dimension table /BIC/DZBWVC_LSDT the rows of table /BIC/DZBWVC_LSDT are qualified.

With the results the join to the fact table /BIC/EZBWVC_LSD is done via access on index /BIC/EZBWVC_LSD~P.

The fact table is sorted by the time dimension. The qualified rows are selected via the index on the fact table. Consequence is that the qualified records are not ordered and not necessarily stored physically together on the same block. The sort order of the fact table cannot be used!

With this access plan the BW query does not benefit from the cluster feature.

The next step is to analyze why the cluster feature is not used.

Step 7a
Check Cluster and Packed Attribute

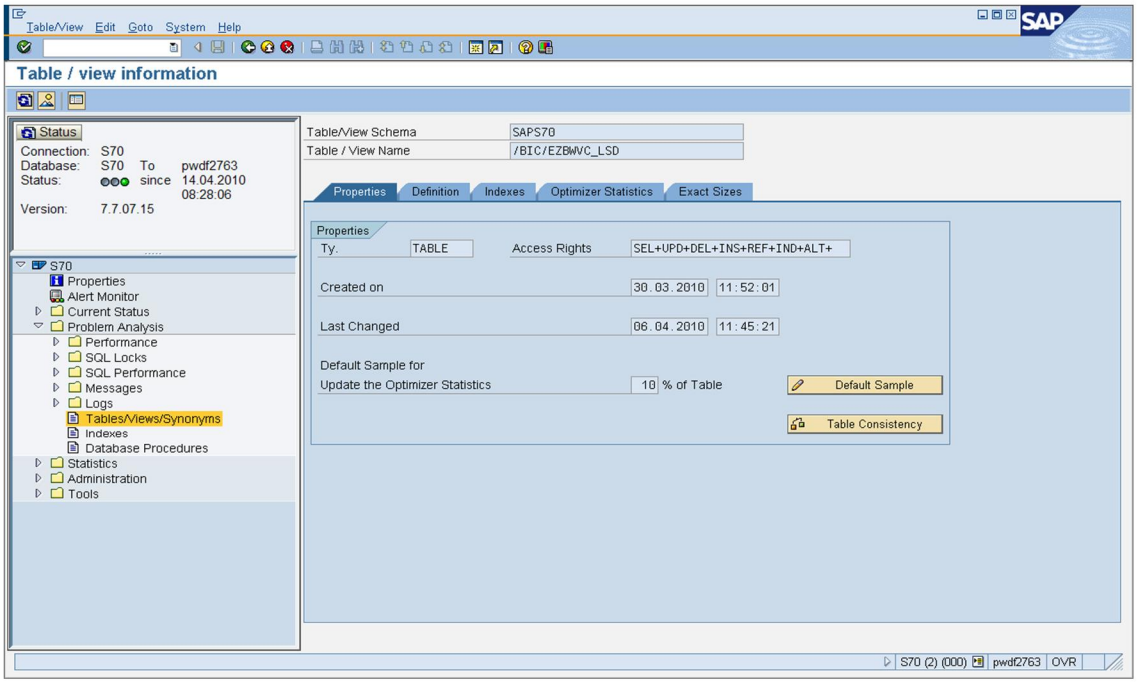

© 2010 SAP AG. All rights reserved. / Page 37

Remember: only fact tables get the attributes CLUSTER and PACKED.

With transaction *DB50 -> Problem Analysis -> Tables/Views/Synonyms* the attributes of fact tables can be checked.

For fact table /BIC/EZBWVC_LSD neither the properties show any information about the CLUSTER or PACKED attribute nor is the tab *Data Storage* shown.

This is a first indicator that the fact table does not have the attributes CLUSTER and PACKED.

Another possibility to check these attributes is to use *DatabaseStudio*.

Step 7b
Check Cluster and Packed Attribute

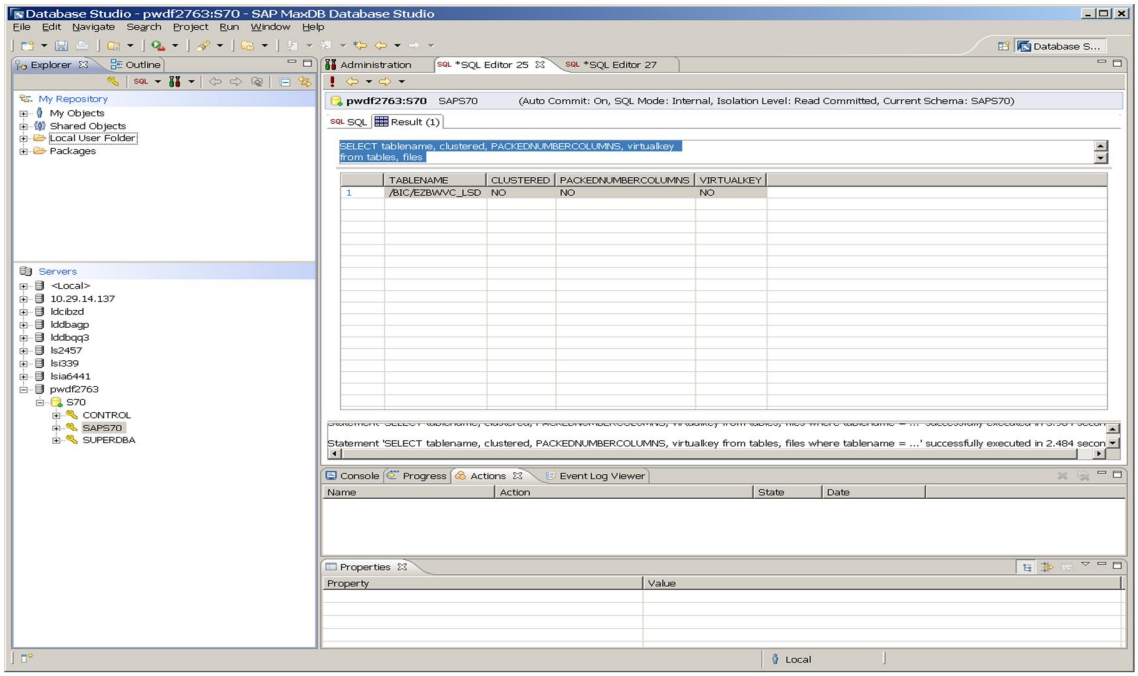

© 2010 SAP AG. All rights reserved. / Page 38

Information about the attributes *CLUSTERED* and PACKED (column: PACKEDNUMBERCOLUMNS*)* is available in the system table *FILES.*

The information if the fact table has a virtual key can be found in system table DOMAIN.TABLES.

```
SELECT TABLENAME, PACKEDNUMBERCOLUMNS, VIRTUALKEY
from tables, files
where tablename = '/BIC/EZBWVC_LSD' and tableid = fileid
```

Table /BIC/EZBWVC_LSD is neither packed nor clustered and has no virtual key.

The next step is to cluster and pack the fact table /BIC/EZBWVC_LSD and to create a virtual key.

Before we cluster and pack the fact table we check the size of the table to be able to compare the size after the table has been packed.

Step 7c
Check Table Size of the Unpacked Table

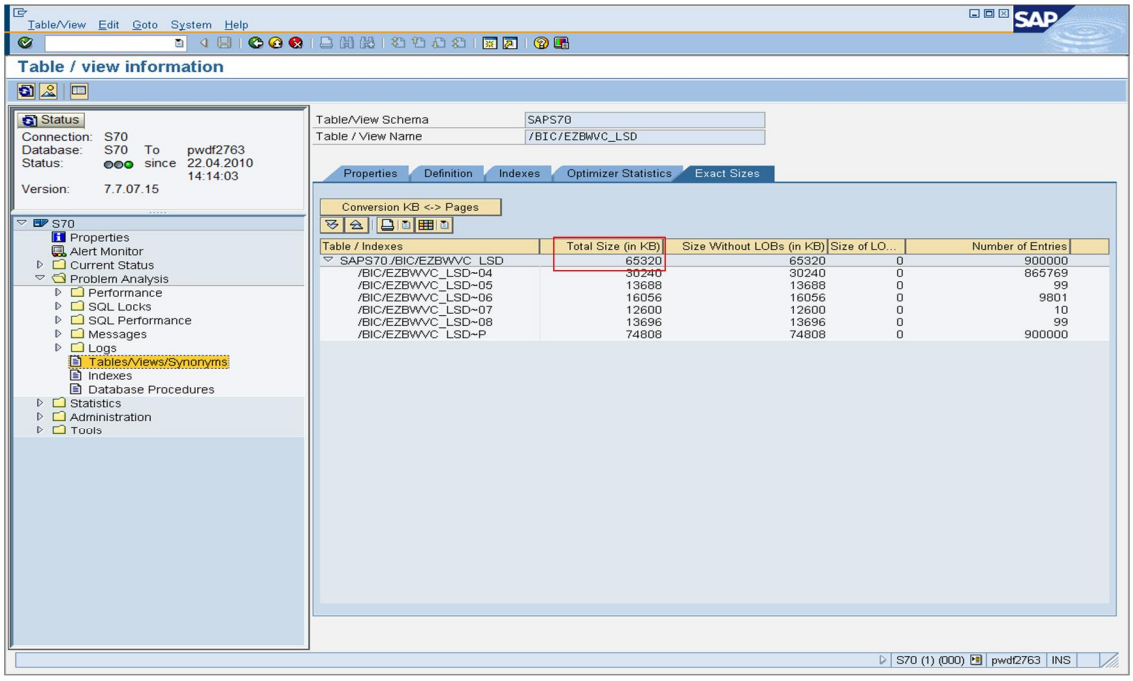

© 2010 SAP AG. All rights reserved.

In transaction DB50 in *Problem Analysis -> Tables/Views/Synonyms* enter the name of the fact table. Use button *Exact Sizes* to get the table size and the number of entries.
The fact table has a size of 65.320 MB.

Remember this size to compare the size after the table has been packed.

Step 8a
Get Cube Name of Fact Table

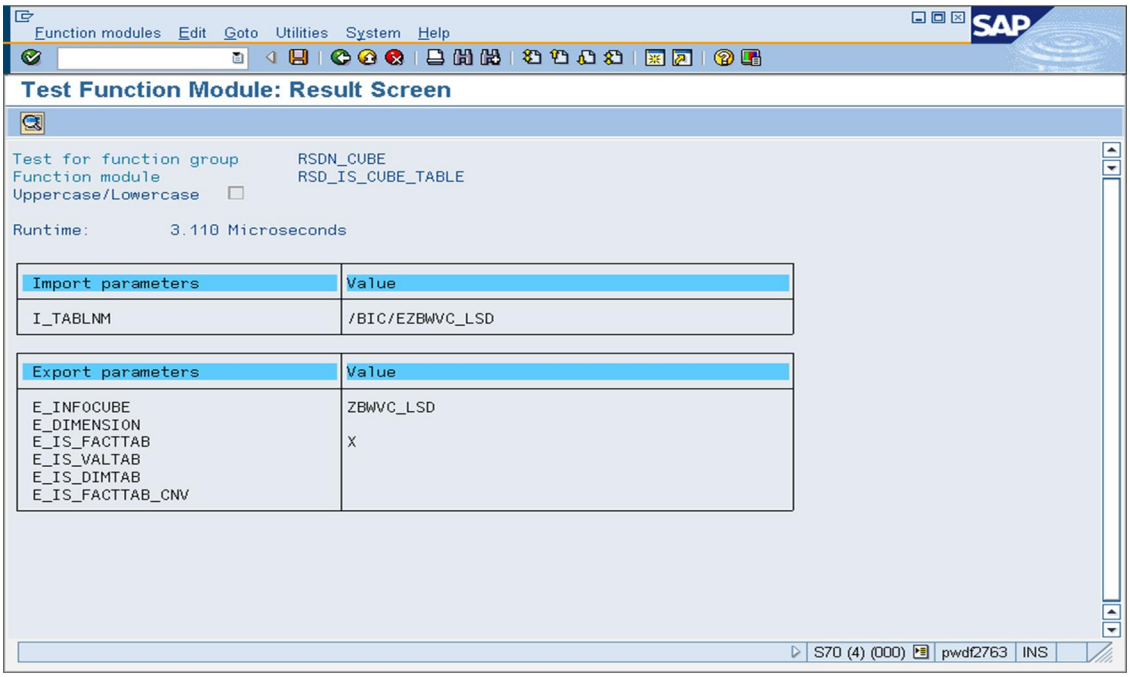

© 2010 SAP AG. All rights reserved. / Page 40

There are 2 possibilities to set the attributes CLUSTERED and PACKED and to create a virtual key.

First possibility:

In a BW system the report RSDU_CLUSTER_FACT_ADA (see note 983845) can be used to add the attributes CLUSTERED/PACKED to the fact tables and to create the virtual key.

Precondition to use that report is that the name of the cube, the fact table belongs to, is known. One can get this information with the function module: RSD_IS_CUBE_TABLE.
The function module RSD_FACTTAB_GET_FOR_CUBE is used to determine the tables that belong to a cube.

During installation/migration of BW-Systems cube and aggregate tables (fact tables) will be created automatically with BW Feature Pack attributes (clustered, packed and virtual key). **This is not implicitly done during an upgrade, therefore a check and execution of the report could be useful.**

All cubes can be converted with the report RSDU_CLUSTER_FACT_ADA .

Step 8b
Set Attributes and Virtual Key

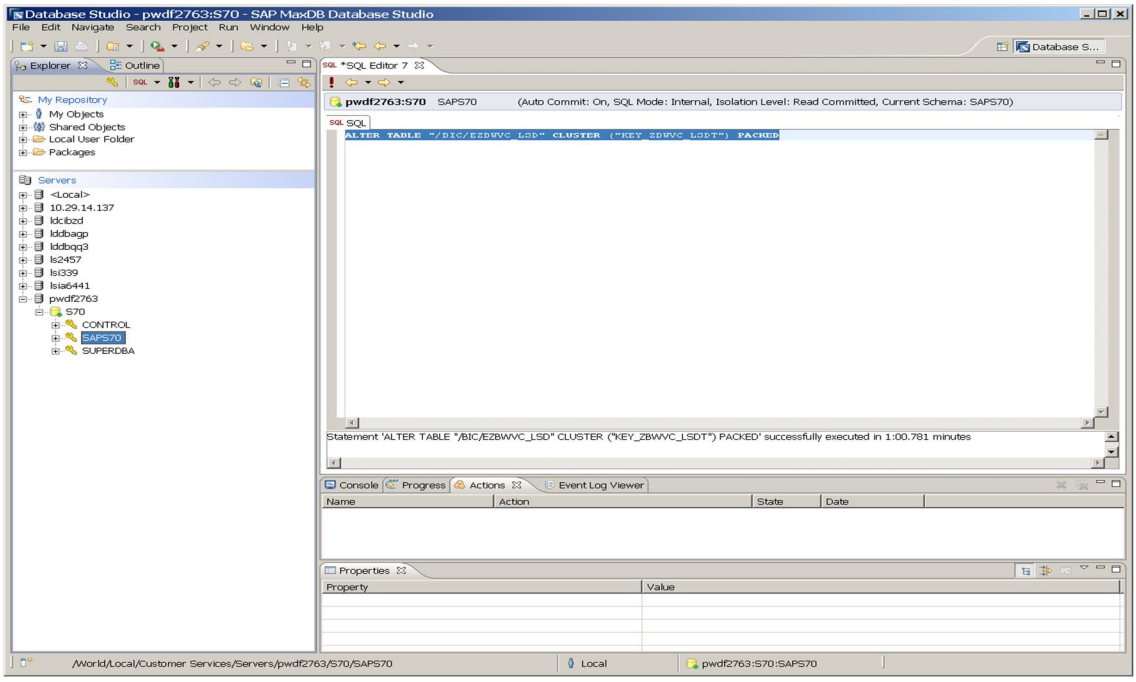

© 2010 SAP AG. All rights reserved. / Page 41

Second possibility:

The cluster/packed attributes and the virtual key can also be set with ALTER TABLE statements via *DatabaseStudio*.

```
ALTER TABLE <tablename> CLUSTER (<virtual-key>) PACKED
```

Example:

```
ALTER TABLE "/BIC/EZBWVC_LSD"
CLUSTER ("KEY_ZBWVC_LSDT") PACKED
```

In the database catalog (system table FILES), the relevant table gets the cluster table attribute. To do this, the system reads the complete table into the cache and stores it in segments on the data volumes during the next savepoint.

Remark: The runtime of the `ALTER TABLE CLUSTER` statement depends on the size of the table. As long as the alter table statement is active with MaxDB Version < 7.7.07.07 the table is locked for all accesses from the application.

As of MaxDB Version 7.7.07.07, the system no longer sets an exclusive lock during `ALTER TABLE CLUSTER`. Read operations can then be carried out in parallel.

Step 8c
Check Cluster and Packed Attribute

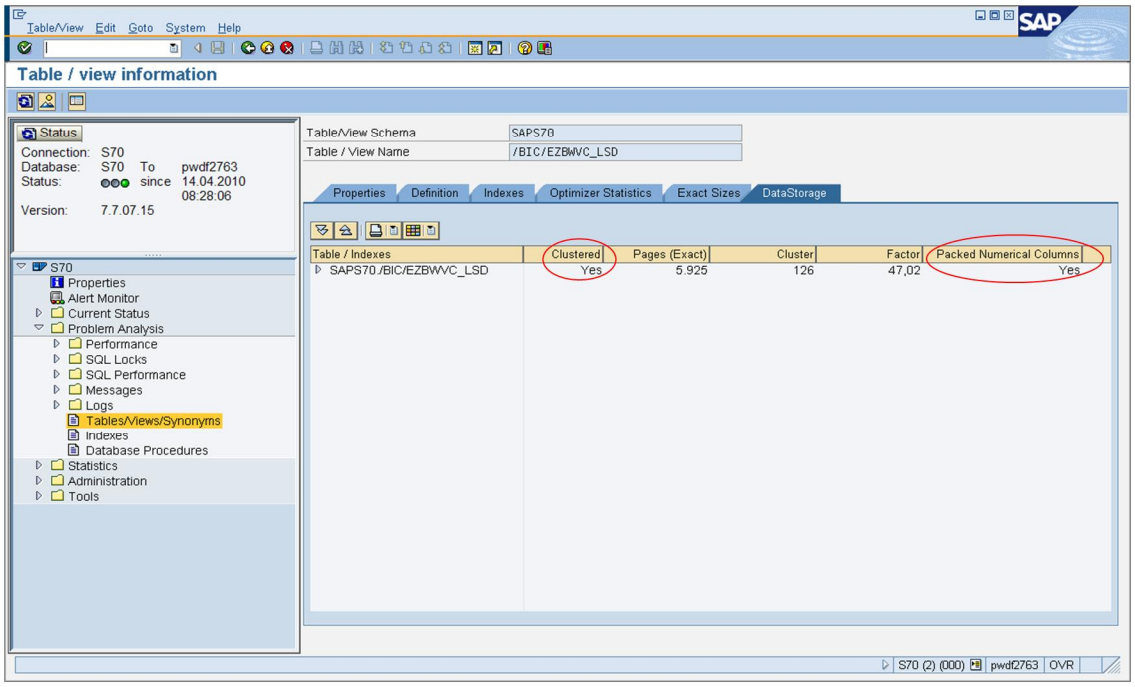

© 2010 SAP AG. All rights reserved. / Page 42

After the ALTER TABLE statement has been executed sucessfully, the additional tab *Data Storage* is displayed in transaction *DB50 –> Problem Analysis –> Tables/Views/Synonyms*.

The CLUSTER and PACKED attributes can be checked here now as well.

Column FACTOR shows that the factor (cluster grade) is good for this table with a value of 47,02.

Notice: When **the FACTOR is close to the value 1 the data is not clustered**. Reason for that could be that there is not enough free space to cluster the data. (see slide *MaxDB Data Storage for BW - Physical Table Clustering – Free Blocks).*

To use the complete functionality of the MaxDB BW Feature Pack including the Star Join Optimization the fact attribute should be checked next.

Step 8d
Check Table Size of the Packed Table

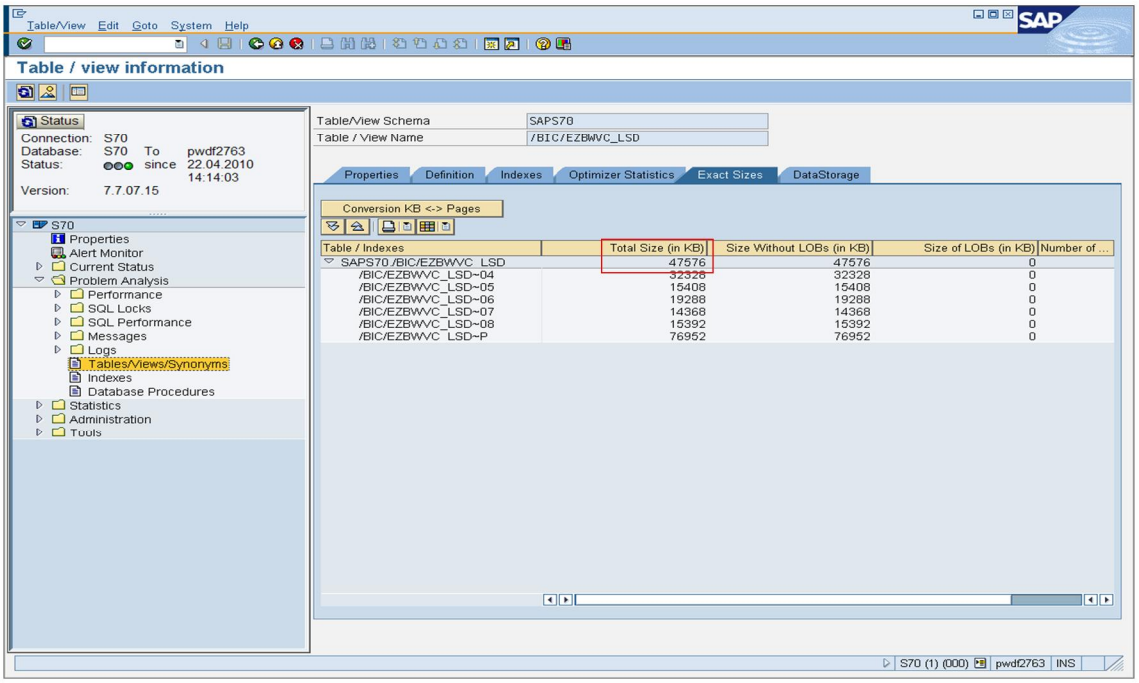

© 2010 SAP AG. All rights reserved. / Page 43

In transaction DB50 in *Problem Analysis -> Tables/Views/Synonyms* enter the name of the fact table. Use button *Exact Sizes* to get the table size and the number of entries.
The fact table has a size of 47.576 MB after it has been packed.

Remember: the size was about 65.3 MB before the table had been packed.

Step 9
Check Attribute Fact

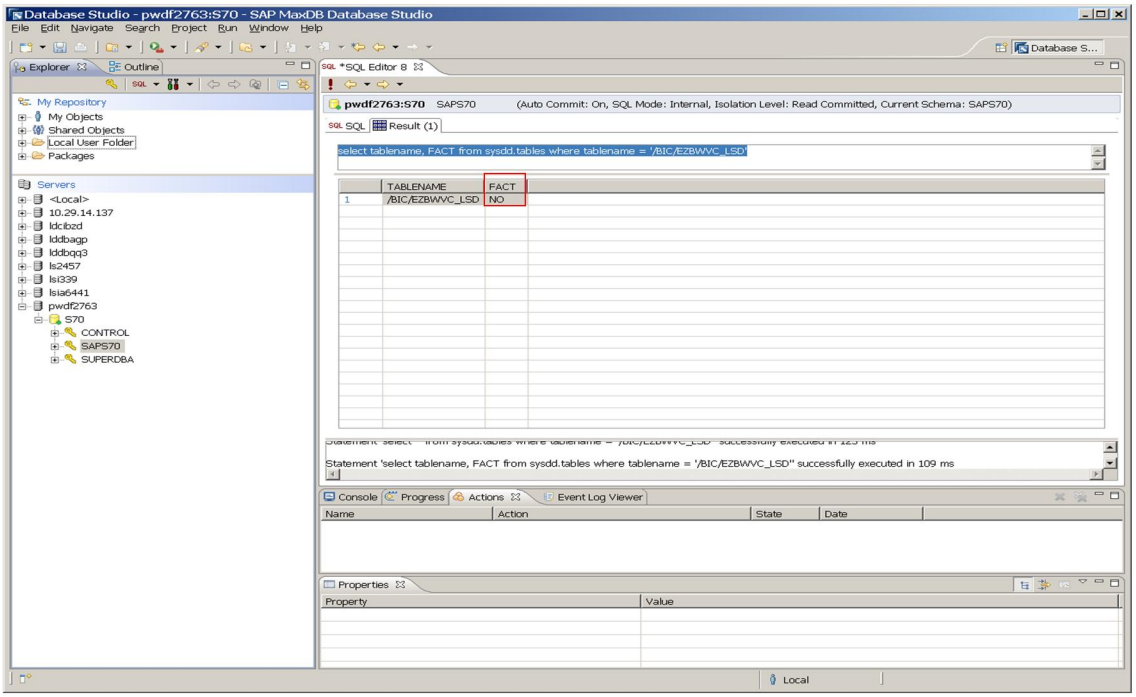

© 2010 SAP AG. All rights reserved. / Page 44

Remember: fact tables have to be marked as fact tables. This enables the special MaxDB Star Join Optimization.

The fact table which has to be analyzed is /BIC/EZBWVC_LSD.

This cannot be checked via transaction DB50. This information is only stored in the system table sysdd.tables.
The following command has to be executed via DatabaseStudio.

```
SELECT tablename, FACT
FROM sysdd.tables WHERE tablename = '/BIC/EZBWVC_LSD'
```

In this example the fact attribute is missing.

Step 10
Set Attribute Fact

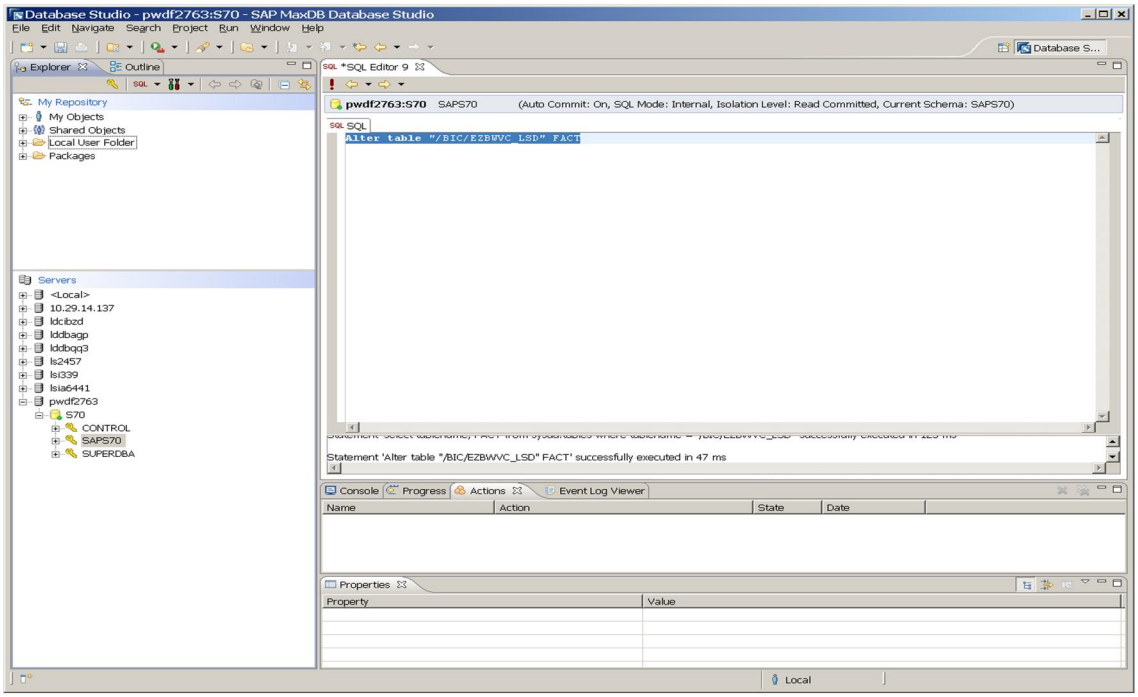

© 2010 SAP AG. All rights reserved. / Page 45

The fact attribute can be set via DatabaseStudio using the following SQL statement:
`ALTER TABLE <tablename> FACT`

Alternatively the Report *RSDU_SET_FACT_ATTR_ADA (*see note 976930) can be executed. This report runs only for a few seconds and ensures that all fact tables are marked.

As an input parameter you can specify a cube name. If you want to set the fact attribute for all fact tables use no input parameter (space).

Step 11
Check Command Monitor

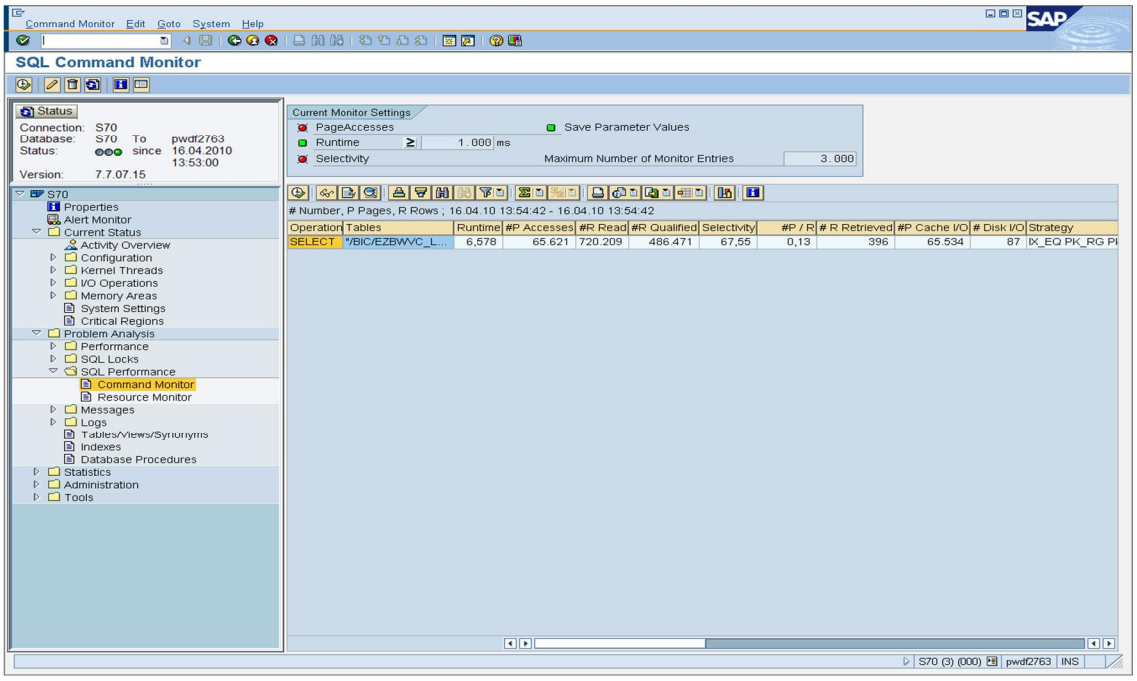

© 2010 SAP AG. All rights reserved. / Page 46

Once the recommended settings are done, run the query again and check the command monitor output.

In our session the database has been restarted to be sure that the cache is empty and that the data has to be read into the cache from disk.

The runtime of the query could be reduced to less than 7 seconds. The number of physical I/O could be reduced to less than 90 I/Os.

What is the new explain plan of this query?

Step 12
Check New Explain

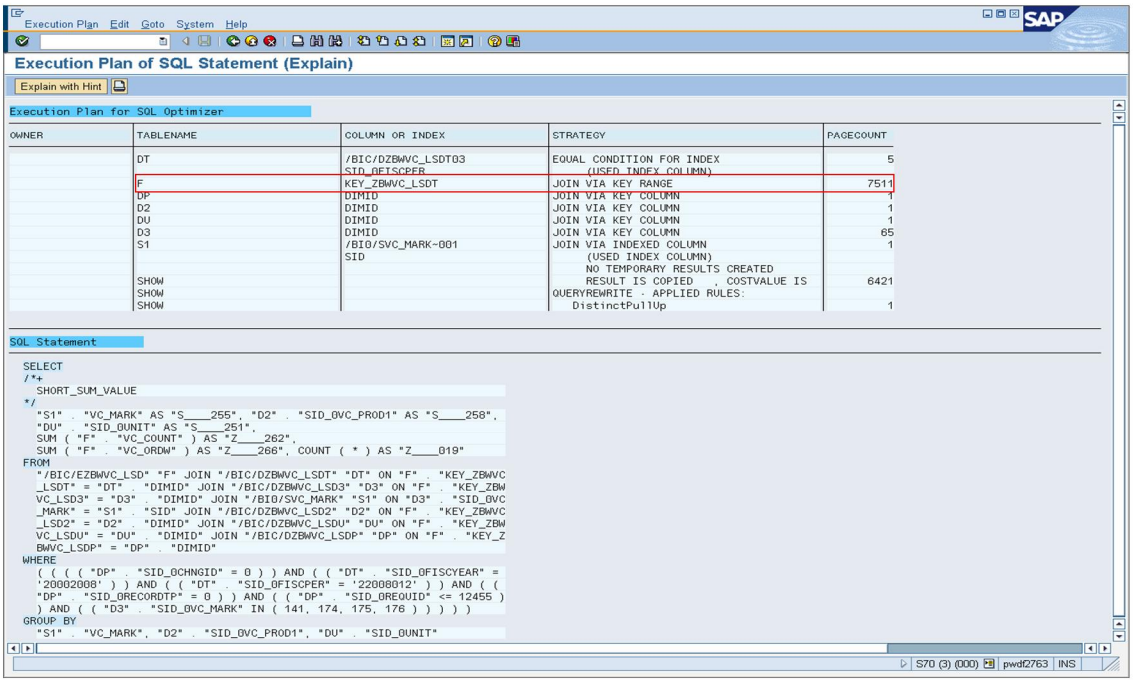

© 2010 SAP AG. All rights reserved. / Page 47

The process order of the tables is the same, but now the (virtual) key access from the time dimension to the fact table is used.

The system now can profit from the physical clustering of the fact table.

JOIN VIA KEY RANGE

Select time dimension via Equal condition for index

Join fact table via (virtual) key range

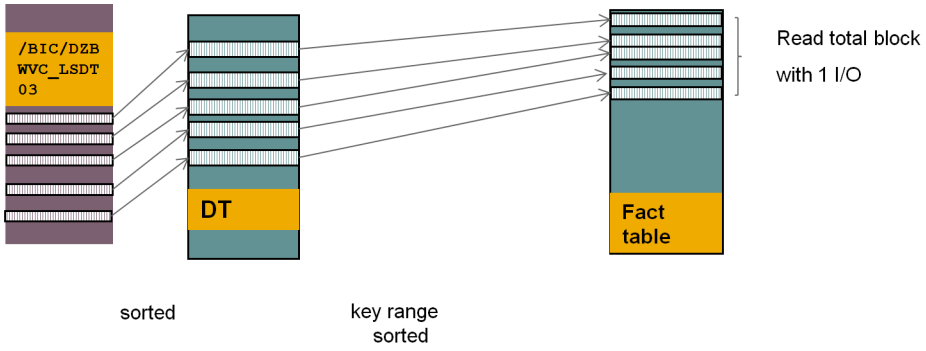

With this access plan the BW Query benefits from the cluster feature

© 2010 SAP AG. All rights reserved. / Page 48

This slide shows the first execution step of the BW-Query.

Starting with the equal access on index /BIC/DZBWVC_LSDT03 of the time dimension table /BIC/DZBWVC_LSDT the rows of this table are qualified. With the results the join to the fact table /BIC/EZBWVC_LSD (sorted**) is directly done** via the (virtual) key access.

The qualified records are ordered and stored physically together on the same block.

With this access plan the BW Query benefits from the cluster feature. In this example several pages can be read with one I/O.

Step 13
Is the Hash Join Used?

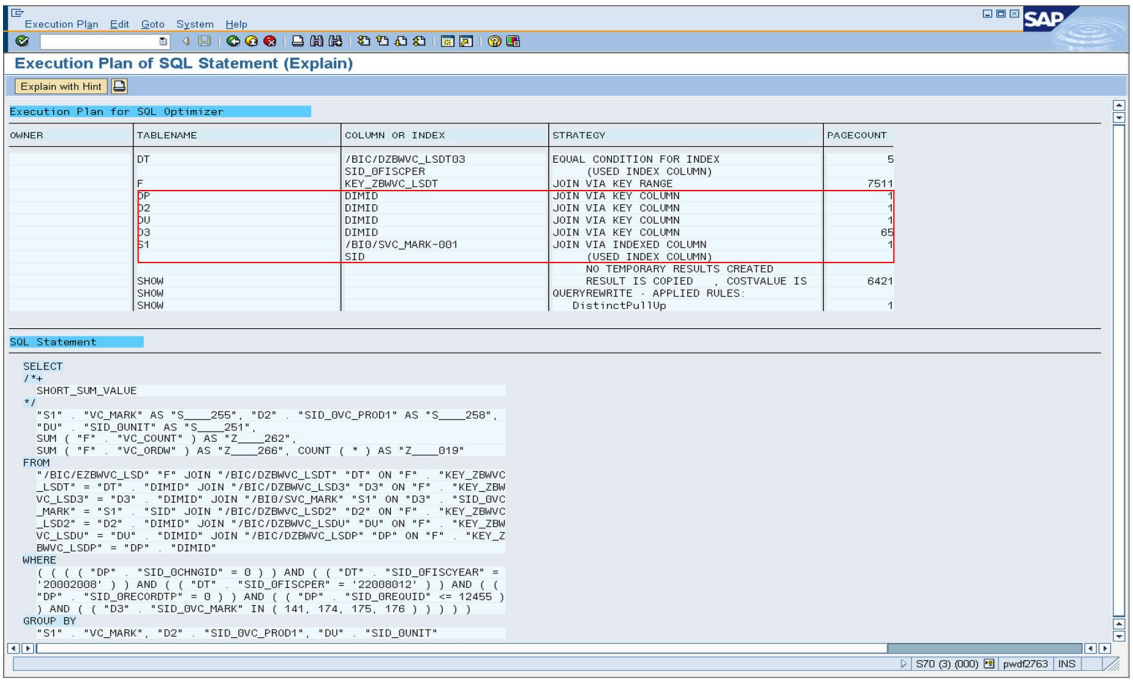

© 2010 SAP AG. All rights reserved. / Page 49

The access plan has already been improved. The statement can be executed faster than in the beginning. However this is still not the best access plan. As you can see the hash join is not used.

Step 14
Check Hash Join Parameter

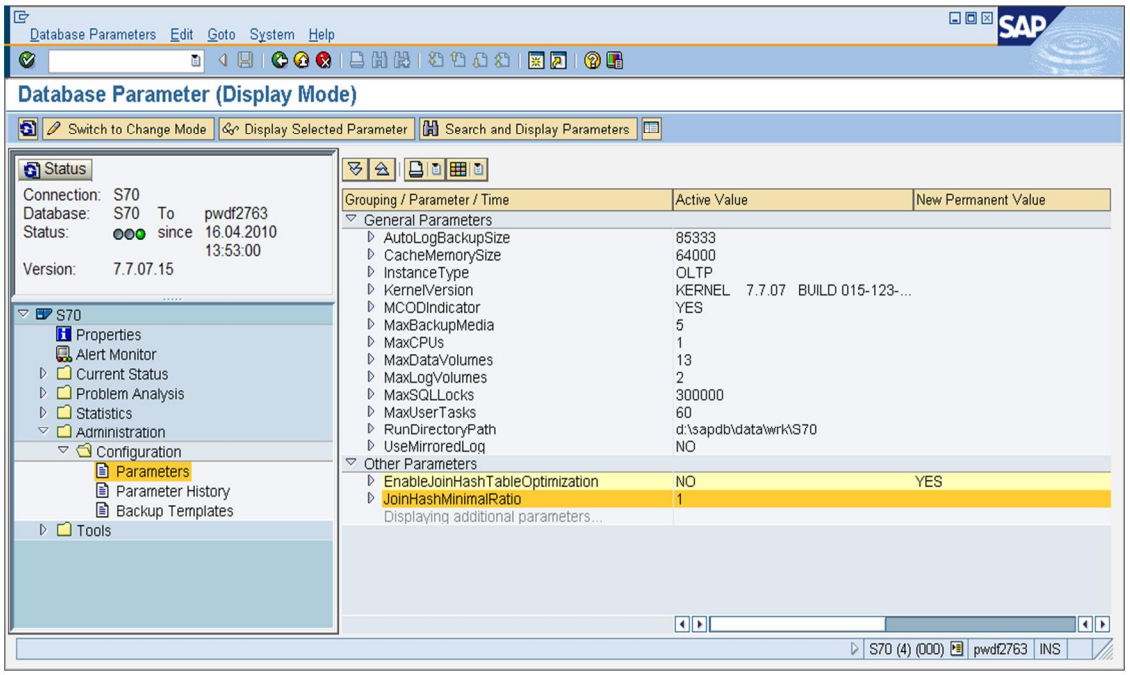

© 2010 SAP AG. All rights reserved. / Page 50

To figure out why the hash join is not used ,check the following parameters:

JoinHashMinimalRatio – is set to 1 = Default = ok

EnableJoinHashTableOptimization –is currently set to NO
The hash join usage is disabled and can therefore not be used.

To activate the hash join the parameter *EnableJoinHashTableOptimization* has to be set to YES. This can be done during online operation in transaction DB50.

Step 15
Check Command Runtime

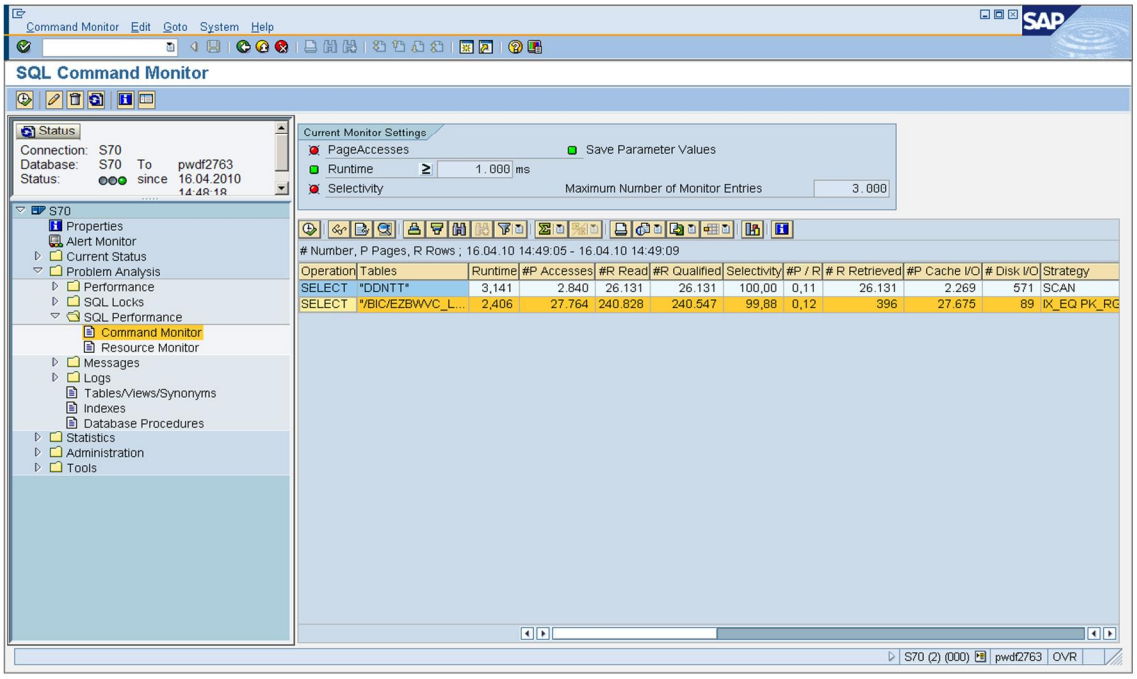

© 2010 SAP AG. All rights reserved. / Page 51

After the Hash Join Optimization has been enabled with parameter *EnableJoinHashTableOptimization,* the runtime could be further reduced to less than 2,5 seconds.

Step 16
Check Explain plan

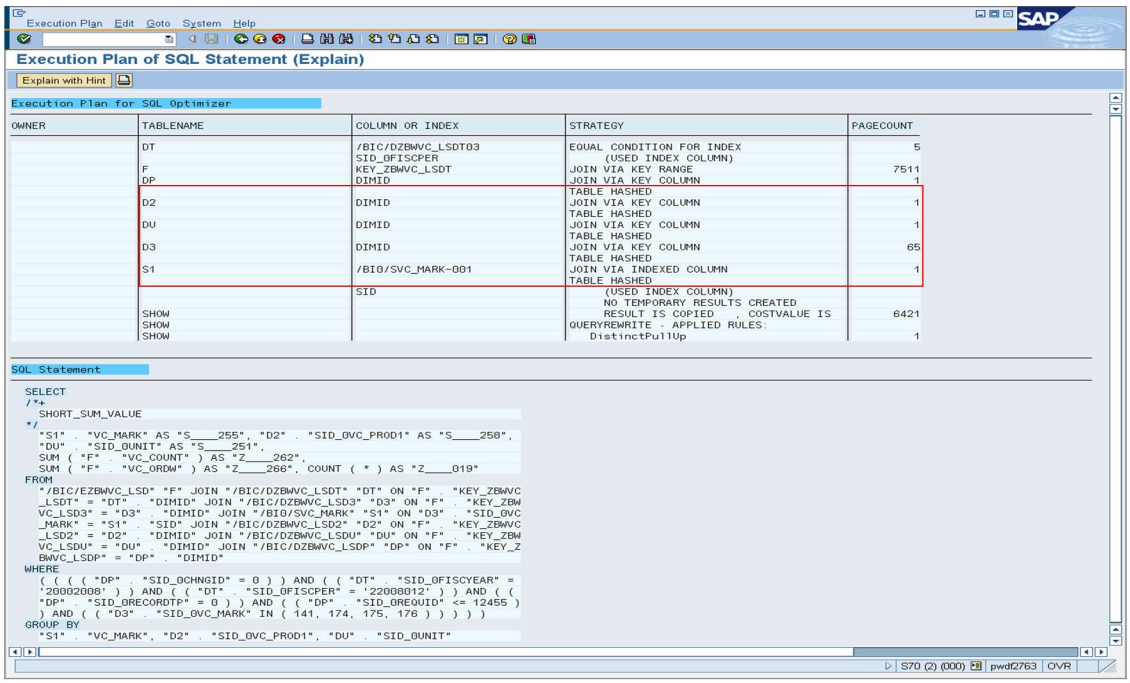

© 2010 SAP AG. All rights reserved. / Page 52

The explain output shows that the hash join is used for the join steps between the dimension tables.

The dimension tables are very small in this example so the runtime benefit is only some seconds.

Notice: This is an example for an optimal used strategy in a BW environment. Most BW queries can be optimized like this one. Nevertheless the shown execution plan must not be the same for all BW queries. In some cases a different strategy might be better.

First access starts with the time dimension (alias DT)

Second access is the fact table – take care that the access is done via key range then the cluster feature is used.

The following dimension tables should be joined via Hash Join.

Agenda

1. Introduction
2. MaxDB Data Storage for BW
   - Physical Table Clustering
   - Table Compression
3. Logical Table Clustering
4. MaxDB Optimizer
   - Optimized Star-Join-Processing
   - Optimized Hash Join
5. Performance Analysis and Monitoring
6. Summary

© 2010 SAP AG. All rights reserved. / Page 53

The target group of this session are customers with SAP Business Warehouse systems running on SAP MaxDB, SAP partners and SAP employees who are doing BW performance analysis on SAP MaxDB systems.

As of SAP MaxDB version 7.6, MaxDB offers special features, which significantly improve the performance of SAP Business Warehouse applications.

This session provides information about these MaxDB features called **SAP MaxDB BW Feature Pack.**

The SAP Business Warehouse application itself is not part of this session.

This session is based on the internal development System S70. It is a SAP NW 7.0 system with limited BW functionality. The MaxDB Version is 7.7.07.

This session consists of a theoretical part and a second part where a performance analysis including checks of the usage of the BW Feature Pack is done directly on the system S70.

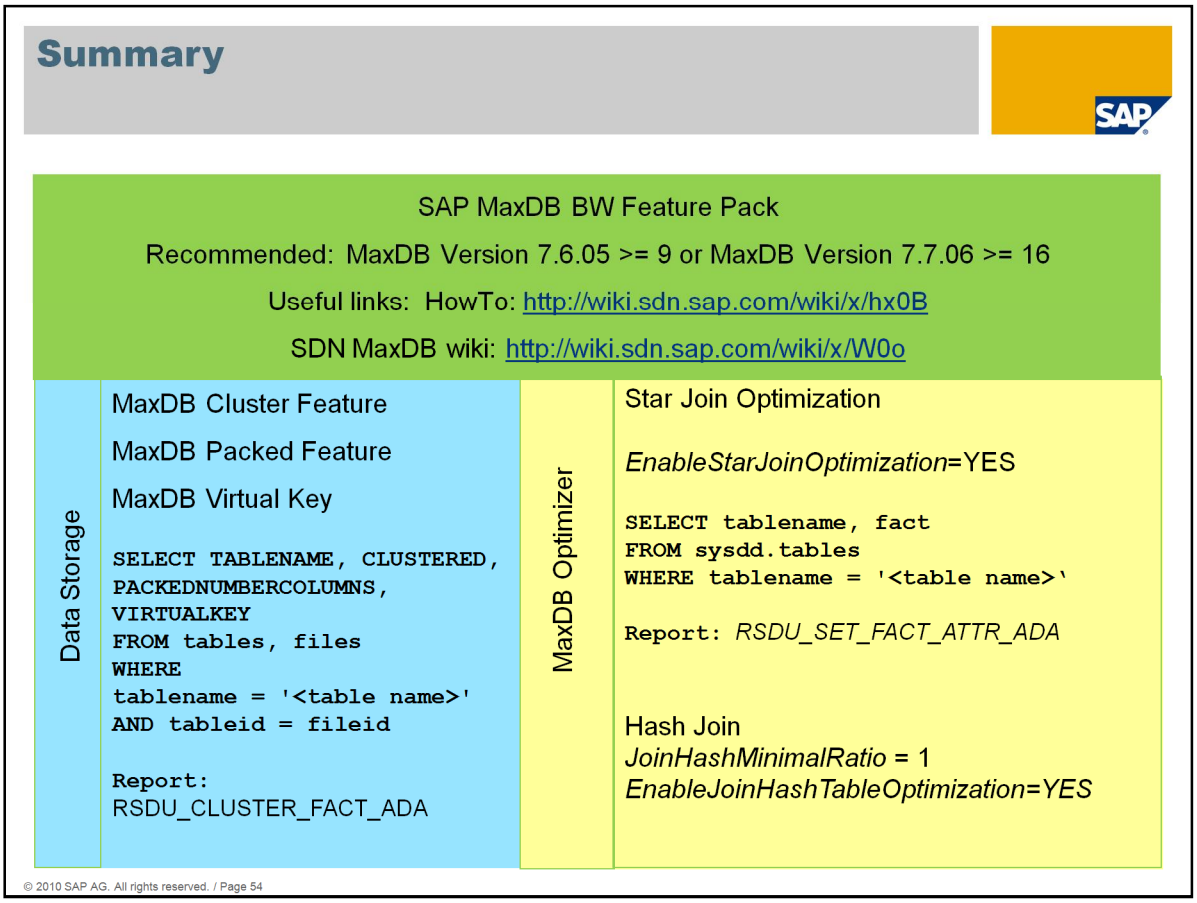

The minimum version which is needed to use the BW feature pack is MaxDB Version 7.6.01.

The Feature Pack is automatically activated during installation/migration, but not during a SAP upgrade. So when there are performance problems in the BW system check if the Feature Pack is activated.

The BW Feature Pack influences the data storage. To check if the data storage related features are activated, select the system tables FILES and TABLES. Run the report RSDU_CLUSTER_FACT_ADA to activate the cluster, packed and virtual key features.

The second part of the BW Feature Pack influences the MaxDB optimizer. Check the parameter setting for the Star Join Optimization. Precondition that the Star Join Optimization can be used is the fact attribute. Check as well that the fact attribute is set correctly for the fact tables.

Precondition that the hash join can be used is the correct parameter setting listed on this slide.

Questions and Answers

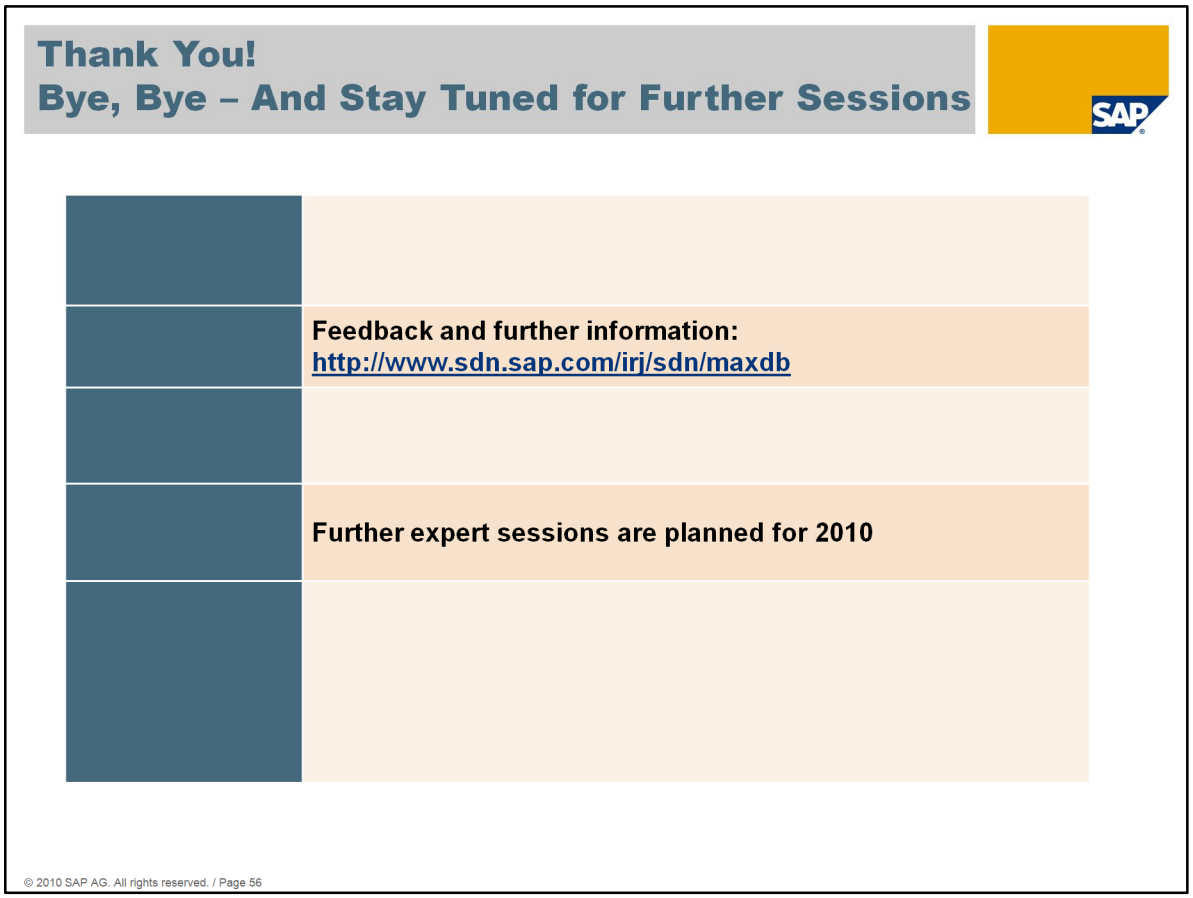

Thanks a lot for your interest in this session.

The recording of this session will be available on maxdb.sap.com/training soon. There you'll find also the recordings of all previously held Expert Sessions about SAP MaxDB.

Further Expert Sessions about SAP MaxDB are currently planned. It will be posted on our SDN page http://www.sdn.sap.com/irj/sdn/maxdb if and when these sessions will take place.

© 2010 SAP AG. All Rights Reserved

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

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

Microsoft, Windows, Excel, Outlook, and PowerPoint are registered trademarks of Microsoft Corporation.

IBM, DB2, DB2 Universal Database, System i, System i5, System p, System p5, System x, System z, System z10, System z9, z10, z9, iSeries, pSeries, xSeries, zSeries, eServer, z/VM, z/OS, i5/OS, S/390, OS/390, OS/400, AS/400, S/390 Parallel Enterprise Server, PowerVM, Power Architecture, POWER6+, POWER6, POWER5+, POWER5, POWER, OpenPower, PowerPC, BatchPipes, BladeCenter, System Storage, GPFS, HACMP, RETAIN, DB2 Connect, RACF, Redbooks, OS/2, Parallel Sysplex, MVS/ESA, AIX, Intelligent Miner, WebSphere, Netfinity, Tivoli and Informix are trademarks or registered trademarks of IBM Corporation.

Linux is the registered trademark of Linus Torvalds in the U.S. and other countries.

Adobe, the Adobe logo, Acrobat, PostScript, and Reader are either trademarks or registered trademarks of Adobe Systems Incorporated in the United States and/or other countries.

Oracle is a registered trademark of Oracle Corporation.

UNIX, X/Open, OSF/1, and Motif are registered trademarks of the Open Group.

Citrix, ICA, Program Neighborhood, MetaFrame, WinFrame, VideoFrame, and MultiWin are trademarks or registered trademarks of Citrix Systems, Inc.

HTML, XML, XHTML and W3C are trademarks or registered trademarks of W3C®, World Wide Web Consortium, Massachusetts Institute of Technology.

Java is a registered trademark of Sun Microsystems, Inc.

JavaScript is a registered trademark of Sun Microsystems, Inc., used under license for technology invented and implemented by Netscape.

SAP, R/3, SAP NetWeaver, Duet, PartnerEdge, ByDesign, Clear Enterprise, SAP BusinessObjects Explorer 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 other countries.

Business Objects and the Business Objects logo, BusinessObjects, Crystal Reports, Crystal Decisions, Web Intelligence, Xcelsius, and other Business Objects products and services mentioned herein as well as their respective logos are trademarks or registered trademarks of SAP France in the United States and in other countries.

All other product and service names mentioned are the trademarks of their respective companies. Data contained in this document serves informational purposes only. National product specifications may vary.

The information in this document is proprietary to SAP. No part of this document may be reproduced, copied, or transmitted in any form or for any purpose without the express prior written permission of SAP AG.

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. Please note that this document is subject to change and may be changed by SAP at any time without notice.

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.

© 2010 SAP AG. All rights reserved. / Page 57