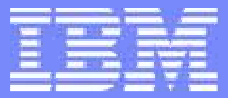

SAP Netweaver BW : Comment augmenter ses performances

Isabelle Claverie-Berge, IT Specialist DB2

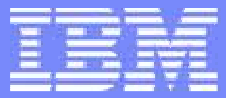

Trademarks

- The following terms are registered trademarks of International Business Machines Corporation in the United States and/or other countries: IBM, AIX, DB2, DB2 Universal Database, Informix, IDS
- SAP, mySAP, SAP NetWeaver, SAP Business Information Warehouse, SAP BW, SAP NetWeaver BI and other SAP products and services mentioned herein are trademarks or registered trademarks of SAP AG in Germany and in several other countries.
- Microsoft, Windows and Windows NT are trademarks of Microsoft Corporation in the United States, other countries, or both.
- UNIX is a registered trademark of The Open Group in the United States and other countries
- Oracle is a registered trademark of Oracle Corporation in the United States and other countries
- LINUX is a registered trademark of Linus Torvalds.
- Other company, product or service names may be trademarks or service marks of others.

Agenda

- DB2 9 , Optimisé pour SAP
  - Partenariat
  - Intégration des produits
  - Innovation technologique
- Les points forts de DB2 et le BI SAP
  - DB2 DPF
  - DB2 MDC
  - DB2 Compression

Architecture SAP et bases de données

- Toutes les applications utilisent la Plateforme d'abstraction fourni par le Serveur d' Application Web SAP
  - Tous les composants SAP NetWeaver sont indépendant de la base de données (à l'exception de SAP BI)
  - Les applications SAP ne se soucient pas de la base de données sur laquelle elles s'appuient

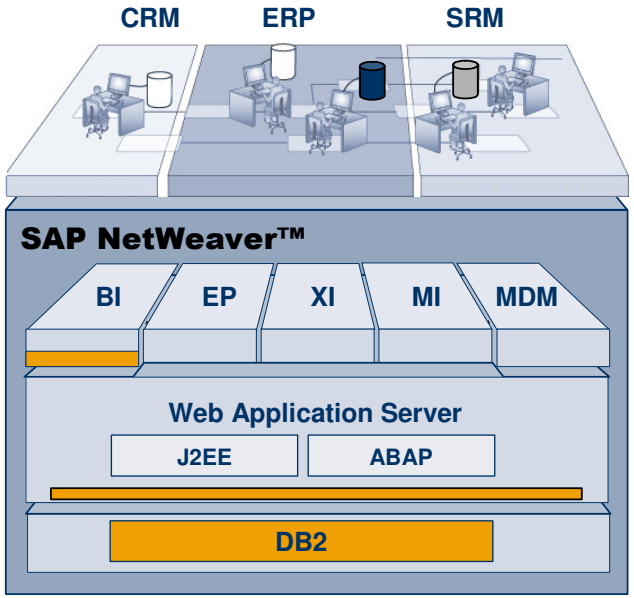

- La base de données optimale pour les clients SAP est celle qui fournit
  - La meilleure integration avec SAP
  - Les meilleures performances
  - La meilleure scalabilité
  - Innovation constante
  - Un produit de qualité
  - Le cout de propriété le plus bas

Ces caractéristiques sont le fil conducteur de l'initiative „optimisé pour SAP"

Les 4 piliers de „DB2 Optimisé pour SAP“

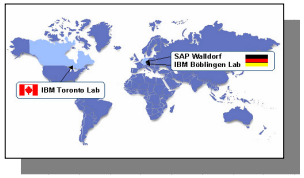

Partenariat

Des équipes communes SAP et IBM qui travaillent ensemble à tous les niveaux

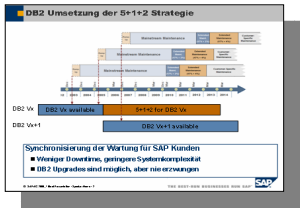

Intégration des produits

Un produit, une stratégie de maintenance, une seule de service

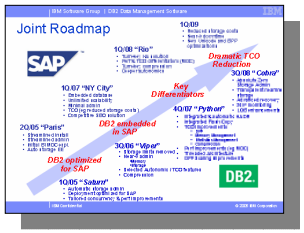

Innovation Technologique

Des plans produits commun jusqu‘à 2008

SAP utilise DB2

SAP est un client DB2 très satisfait

Ce type de partenariat stratégique est un unique entre SAP et IBM

Partenariat SAP-IBM et intégration des produits

- SAP sur DB2 est un projet commun depuis 1993:
  - Des équipes commune de développement SAP+IBM à Walldorf
  - Un Centre d'Integration SAP - IBM à Toronto

    Coopération extrêment proche avec le Développement DB2 à Toronto

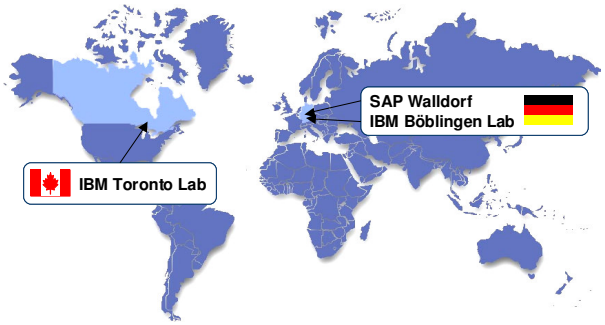

- SAP sur DB2 est un produit complétement intégré

La Stratégie de Maintenance SAP 5-1-2

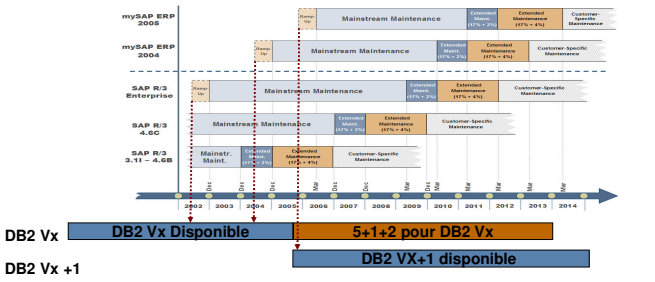

Une valeur unique pour les clients SAP:

- Les mises à jour de DB2 ne sont jamais obligatoires, car la release de DB2 associé au produit SAP reste supportée pendant tout le cycle de vie du système SAP
- Mais on peut aussi choisir de faire évoluer la version de DB2 pour bénéficier des innovations technologiques

THE BEST-RUN BUSINESSES RUN SAP

SAP - IBM support

SAP et IBM maintiennent une structure de support par niveau pour les clients SAP sur DB2 UDB

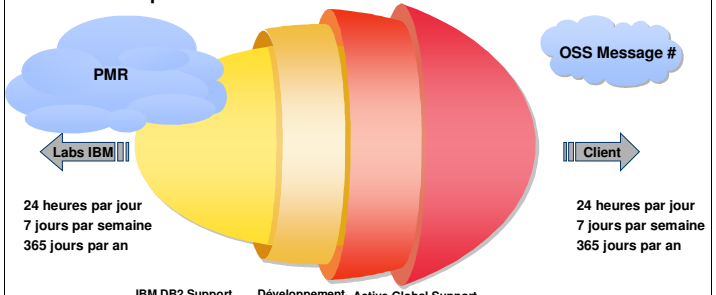

THE BEST-RUN BUSINESSES RUN SAP

Innovation Technologique

- Planning joint des releases SAP+IBM pour DB2 jusqu'en 2008

- Status des produits DB2 „Optimisés pour SAP“
  - DB2 Saturn V8.2.2 disponible depuis Avril 2005
  - DB2 9 disponible depuis Aout 2006
  - DB2 Python planning complet, livraison planifiée le 4ieme trimestre 2007
  - DB2 Cobra planning démarré en Avril 2006

- Record de performance pour DB2 comme le prouve les résultats obtenus lors des benchmarks standard SAP :
  - SD 3-tier benchmark: 168.300 SD User (certificate 2005021)
  - SAP BI 3-tier benchmark (certificate 2005043)

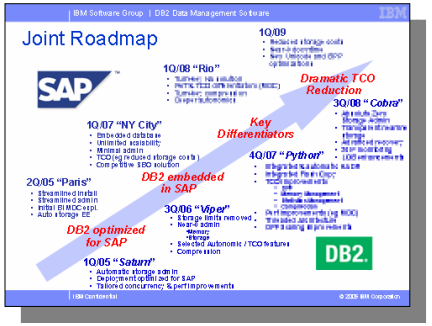

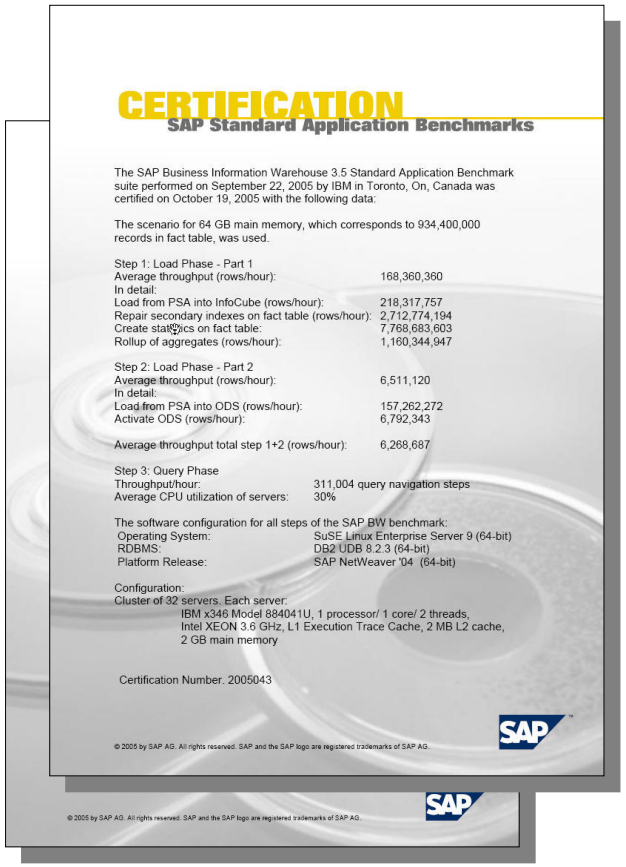

CERTIFICATION

SAP Standard Application Benchmarks

The SAP Business Information Warehouse 3.5 Standard Application Benchmark suite performed on September 22, 2005 by IBM in Toronto, On, Canada was certified on October 19, 2005 with the following data:

The scenario for 64 GB main memory, which corresponds to 934,400,000 records in fact table, was used.

Step 1: Load Phase - Part 1
Average throughput (rows/hour): 168,360,360
In detail:
Load from PSA into InfoCube (rows/hour): 218,317,757
Repair secondary indexes on fact table (rows/hour): 2,712,774,194
Create statistics on fact table: 7,768,683,603
Rollup of aggregates (rows/hour): 1,160,344,947

Step 2: Load Phase - Part 2
Average throughput (rows/hour): 6,511,120
In detail:
Load from PSA into ODS (rows/hour): 157,262,272
Activate ODS (rows/hour): 6,792,343

Average throughput total step 1+2 (rows/hour): 6,268,687

Step 3: Query Phase
Throughput/hour: 311,004 query navigation steps
Average CPU utilization of servers: 30%

The software configuration for all steps of the SAP BW benchmark:
Operating System: SuSE Linux Enterprise Server 9 (64-bit)
RDBMS: DB2 UDB 8.2.3 (64-bit)
Platform Release: SAP NetWeaver '04 (64-bit)

Configuration:
Cluster of 32 servers. Each server:
IBM x346 Model 884041U, 1 processor/ 1 core/ 2 threads,
Intel XEON 3.6 GHz, L1 Execution Trace Cache, 2 MB L2 cache,
2 GB main memory

Certification Number. 2005043

SAP

SAP BW and IBM DB2

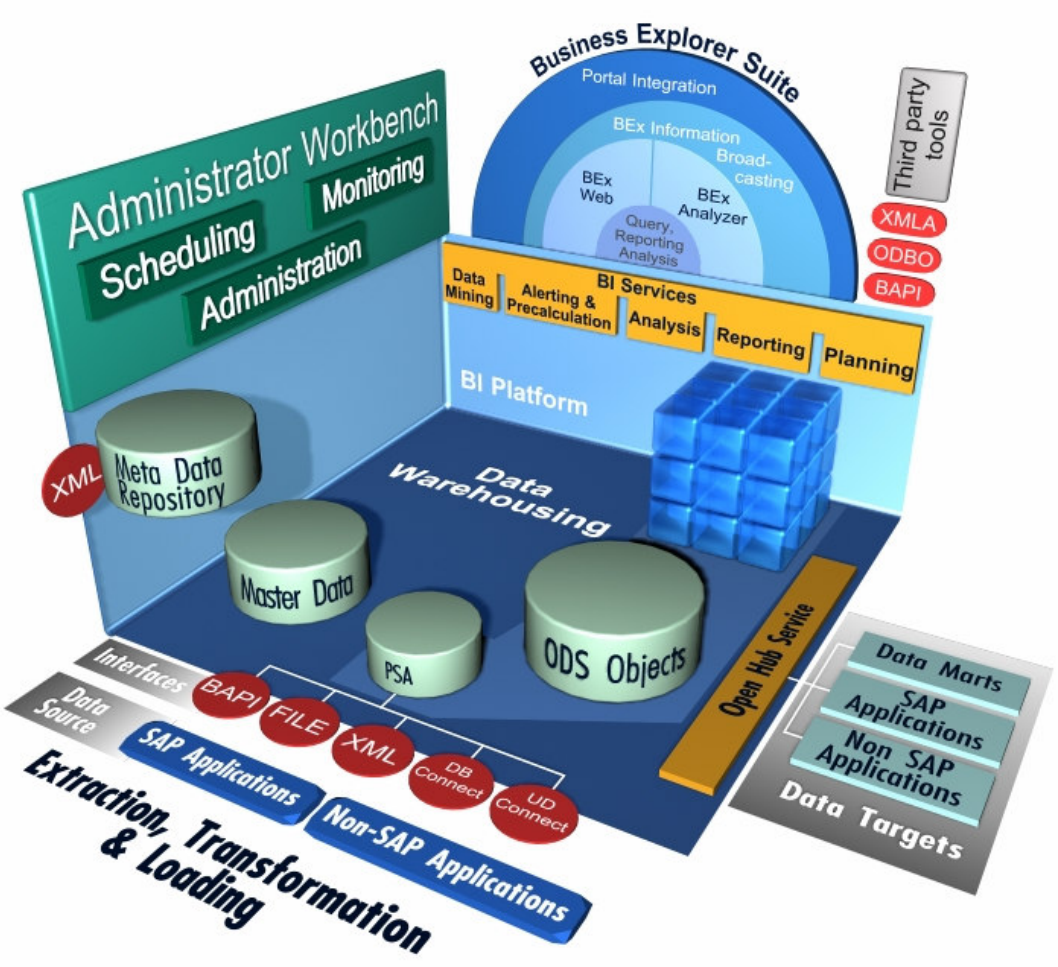

SAP BW architecture uses database functions

- Data load
- Query execution
- Database objects
- Warehouse management
- Analysis/Tuning
- Scalability

SAP BW platform integration

- Transparent
- Database specific

IBM DB2 functions supporting SAP BI

- Optimizer
- Scalability
- Performance
- OLAP Technology

Agenda

- DB2 9 , Optimisé pour SAP
  - Partenariat
  - Intégration des produits
  - Innovation technologique
- Les points forts de DB2 et le BI SAP
  - DB2 DPF
  - DB2 MDC
  - DB2 Compression

IBM DB2 Multi-Partition Concept

- **One database can reside on several separate computers**
  - **Shared nothing (function shipping)**
    - **Each Partition accesses only its local Data**

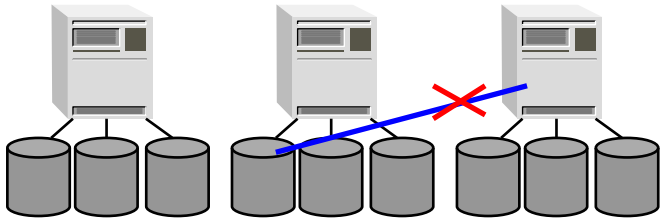

- **Several logical partitions can be on the same machine**
  - **Physical or Logical Partitioning is transparent to the database.**

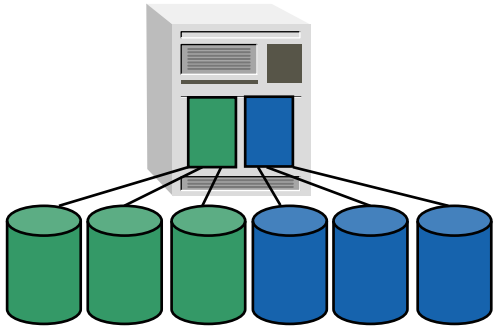

- **Database Catalog on partition 0, DB catalog cache on the other partitions**
- **Fast communication needed**

  (Gigabit Ethernet, Switch)

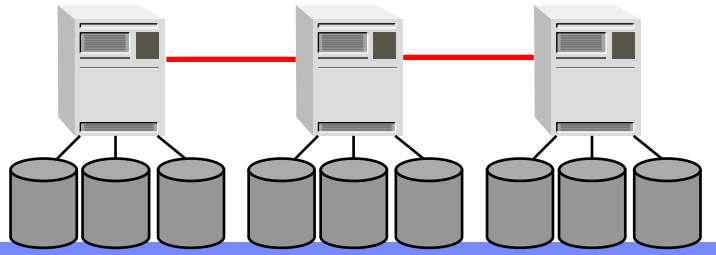

IBM DB2 Parallel Processing

- **Two types of Parallelism**
  - Intra-partition parallelism
    => parallel processing within one partition
  - Inter-partition parallelism
    => operations are executed in parallel on each database partition
- **Scalabilité**: SQL Query performance proportional to number of partitions (BW environment)
- **Tous les ordres SQL**
  UPDATE, DELETE, INSERT, JOINS, GROUP BY, INDEX/TABLE SCANS, SORT
- **Tools**: INDEX Creation, Backup and Restore, Table Reorganization

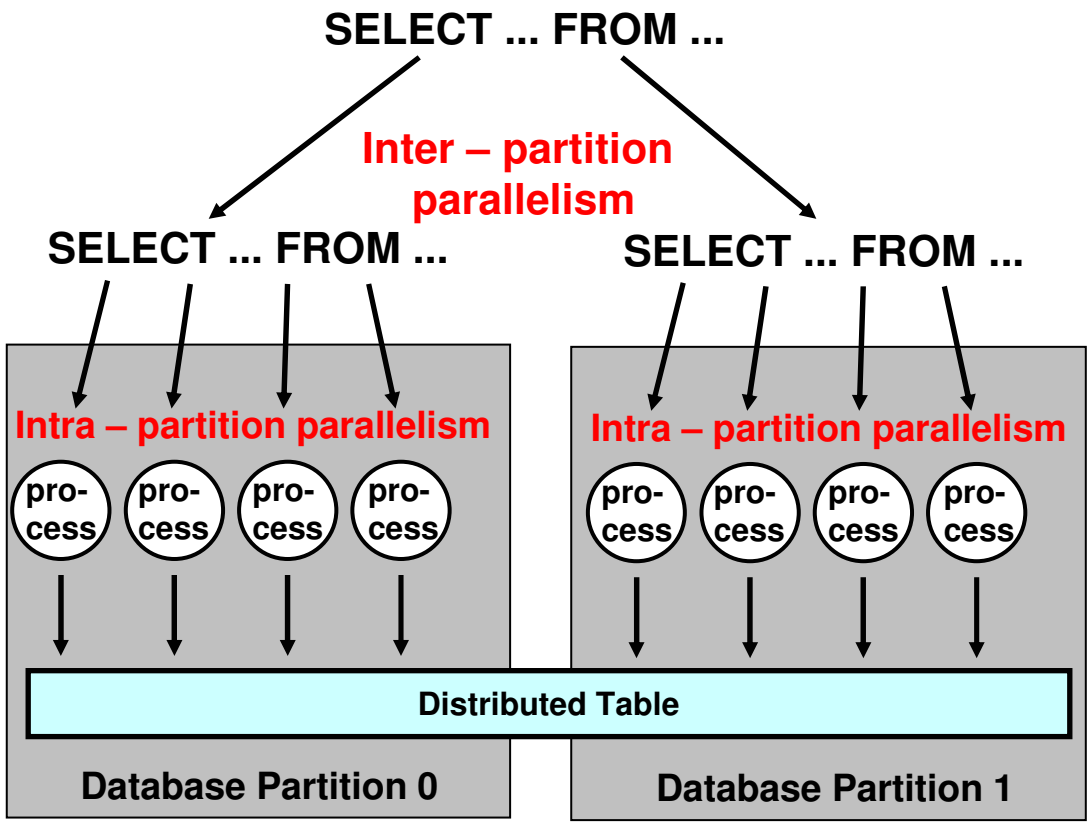

SAP NetWeaver BI on DB2 LUW DPF – Database Layout

- BI Application Servers connected to DB partition 0
  - Fast access to SAP Basis, master data and dimension tables
- Large SAP NetWeaver BI tablespaces distributed over several DB partitions
  - Partitioning Key generated automatically
  - Uniform data distribution
- -> Near linear scalability

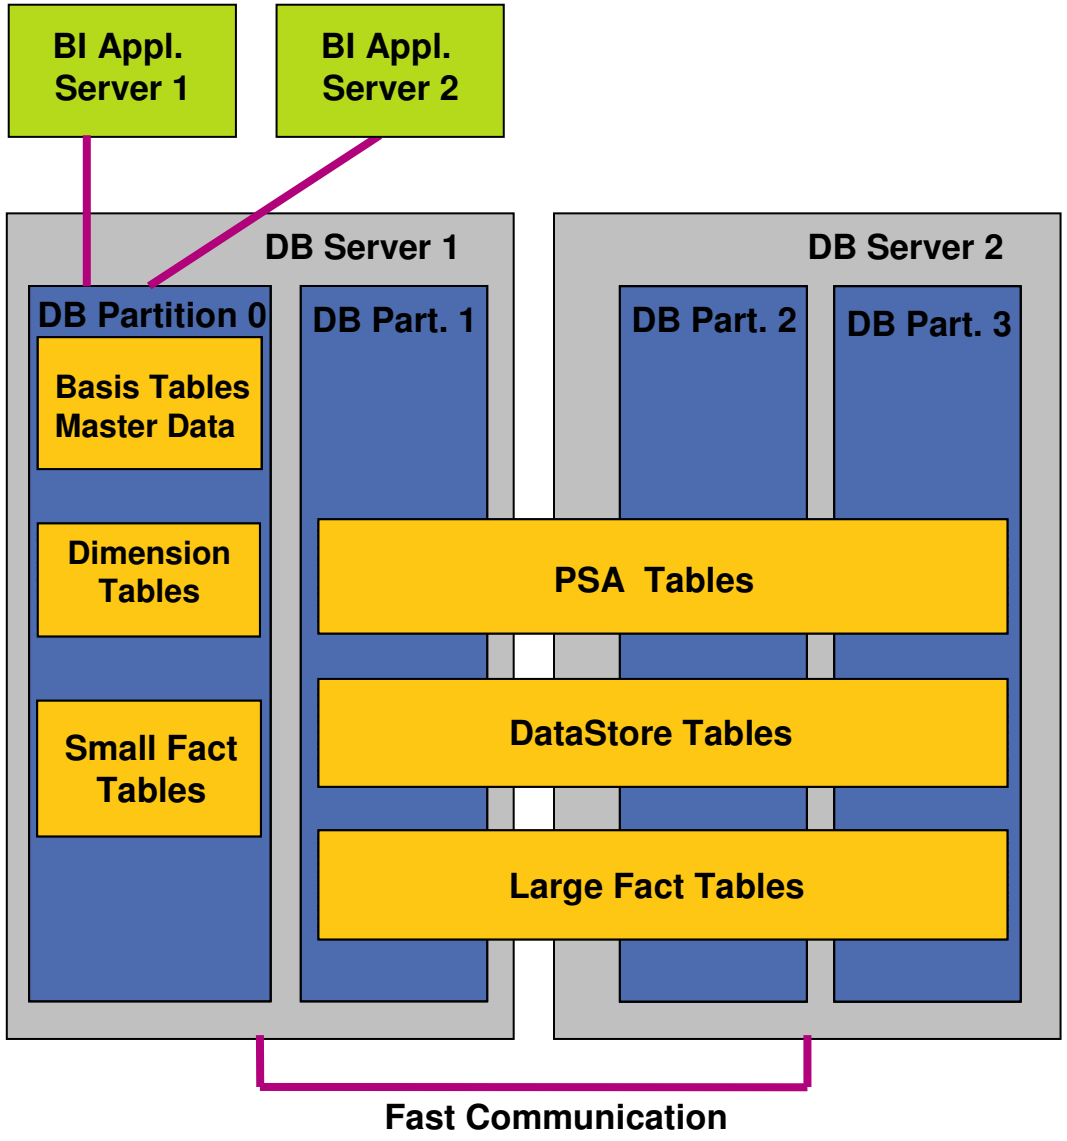

DB2's Shared Nothing Architecture

DB2's Parallel Features for SAP BW, APO, SEM

- Partitions/nodes can be logical or physical
- Configurations can consists of a mixture of logical and physical nodes
- Each node has only access to its own data
- Parallel administration operation of each node
  - Backup & restore
  - Reorganization
  - Update statistics
- The partitioning is transparent to the database application

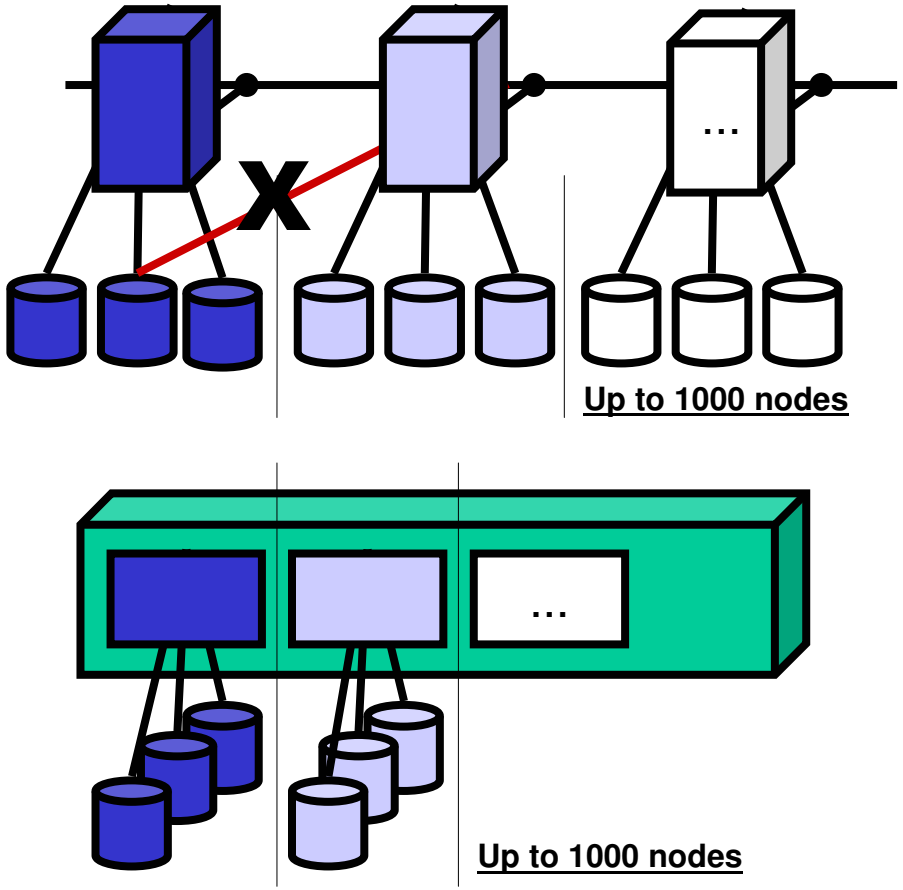

SAP BW – IBM DB2 Specific Features (1)

Features used for SAP BW Query Processing

- Common Table Expression (CTE) = Dynamic Database Views
- DB2 Reduced Optimization: Prepare time (= time to create SQL access plan) of BW queries are reduced by factors
- Multidimensional Star-Joins with dynamic bitmaps; used for queries with multiple restrictions on dimensions
- Hash Join is often used
- Fetch First N Rows / Optimize for N Rows
- Multi-Dimensional Clustering (MDC)

Features used for SAP BW Warehouse Management

- Flexible Database Layout:
  - Multiple Tablespaces with different Page Sizes
  - Multiple Buffer Pools
- Parallel Database: Multiple Database Partitions are used
- Index Clustering on Fact Tables
- Insert Buffering: Faster BW Data Load in Multi-Partition environment

DB2 Dynamic Bitmapped Indexing

Performance

Star Join Example

```
Select ...
  from store st, sales sa,
          period pe, product pr
  where st.store_id = sa.store_id
    and  pr.product_id = sa.product_id
    and  pe.period_desc =sa.period_desc
    and  pe.year between 1995 and 1996
    and  pr.producer in ('TETRA','WISA')
  group by st.store_id, pe.year
```

1. Do Semijoins
2. Generate bitmaps
3. AND bitmaps
4. Select matching records

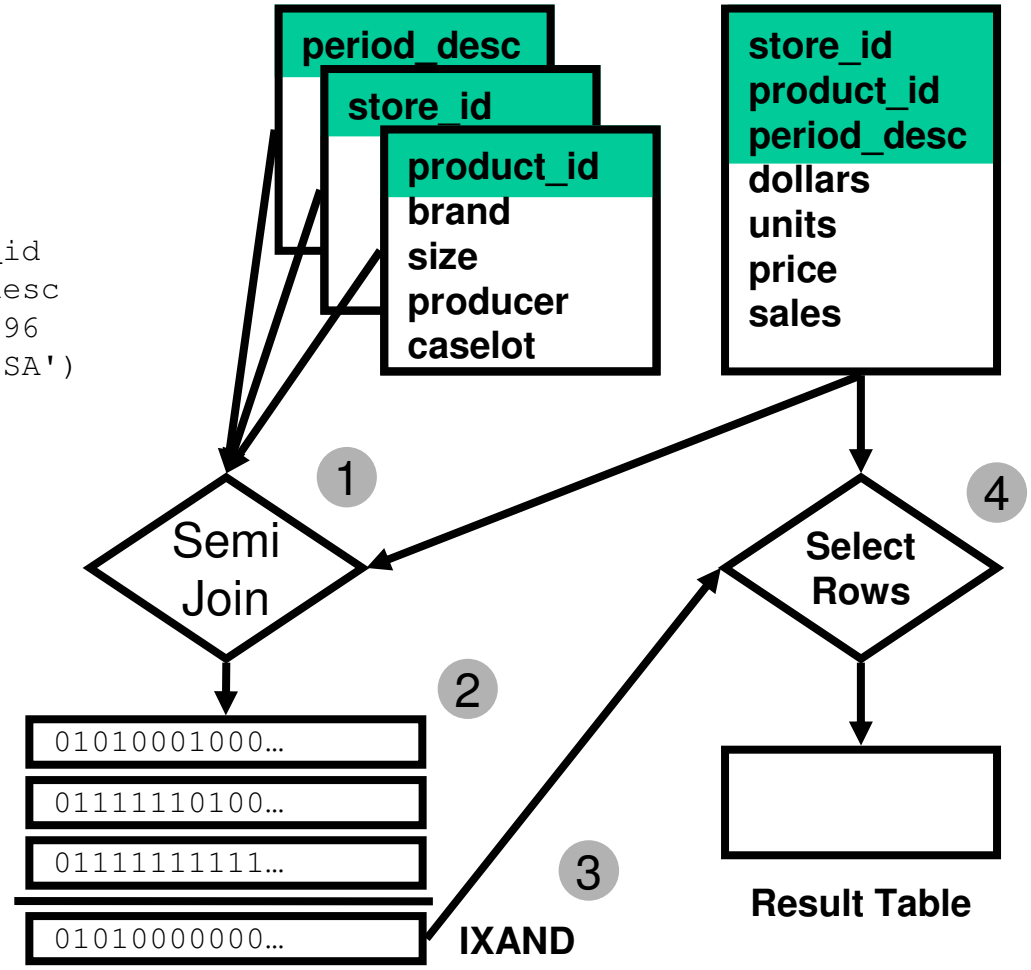

SAP BW – IBM DB2 Specific Features (2)

Tailored Database Statistics

- SAP Basis tables: DB2 Basic Statistics
- SAP BW ODS, PSA tables: DB2 Basic Statistics
- SAP BW Fact, Dimension, Master Data tables: DB2 Distribution Statistics

DB2 Features used in all SAP Products

- Online Backup
- Online Table Reorganization
- Incremental Backup
- DB6COCKPIT

Proof Points: DB2 SAP Performance

Company, Germany: SAP BW Proof-of-Concept

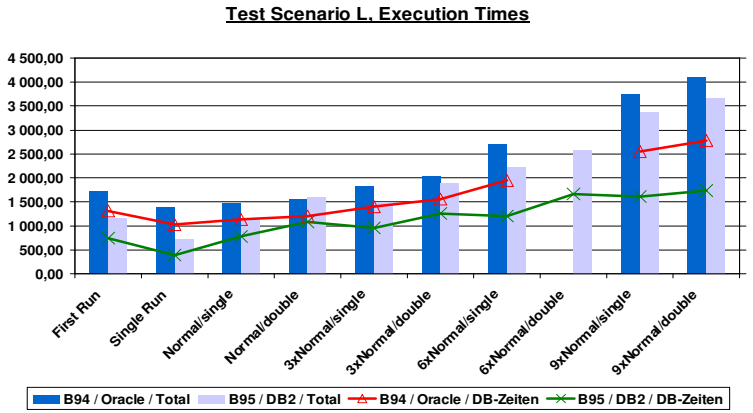

- Overall performance with DB2 better in 8 out of 9 test cases
- DB2 better in all cases if looking at pure DBMS times

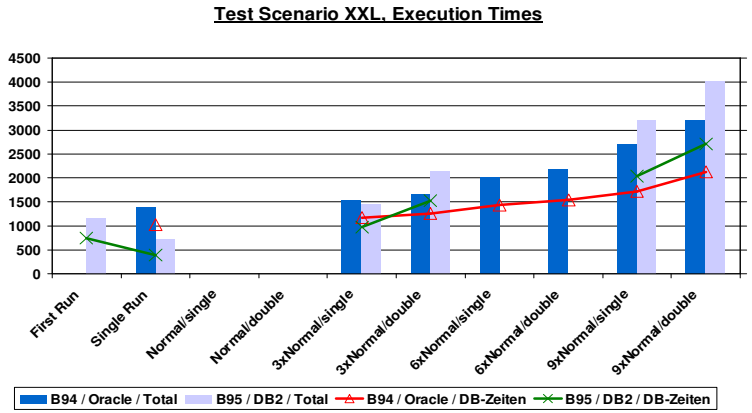

- Overall performance with Oracle slightly better than with DB2
- Extremely little time allowed by customer for DB2 tuning, while Oracle environment was based on existing production system (tuned since years) !

Agenda

- DB2 9 , Optimisé pour SAP
  - Partenariat
  - Intégration des produits
  - Innovation technologique
- Les points forts de DB2 et le BI SAP
  - DB2 DPF
  - DB2 MDC
  - DB2 Compression

MDC Overview (I)

- MDC organizes table data along one or more columns defined as MDC dimensions
- Records with the same values in the dimension columns are stored physically together in blocks of pages
- Blocks correspond to tablespace extents
- MDC block indexes point to blocks instead of single rows
- If more than one MDC dimension is selected a compound block index might be created additionally

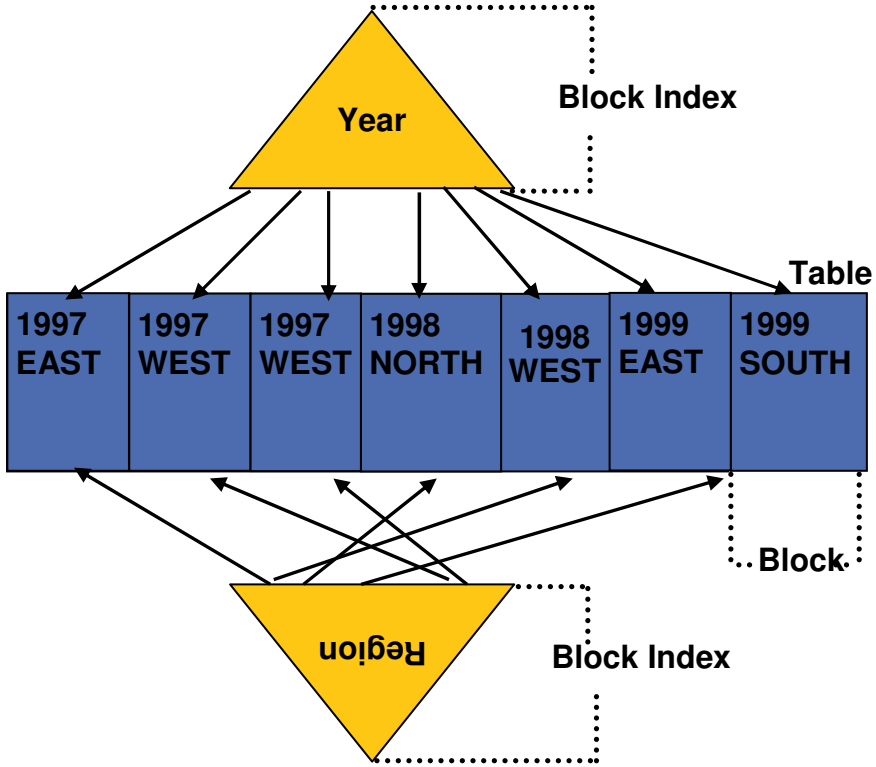


```
CREATE TABLE MDCTAB
  (YEAR CHAR(4),
   REGION VARCHAR(5), …)
  ORGANIZE BY DIMENSIONS(YEAR, REGION)
```

MDC Overview (II)

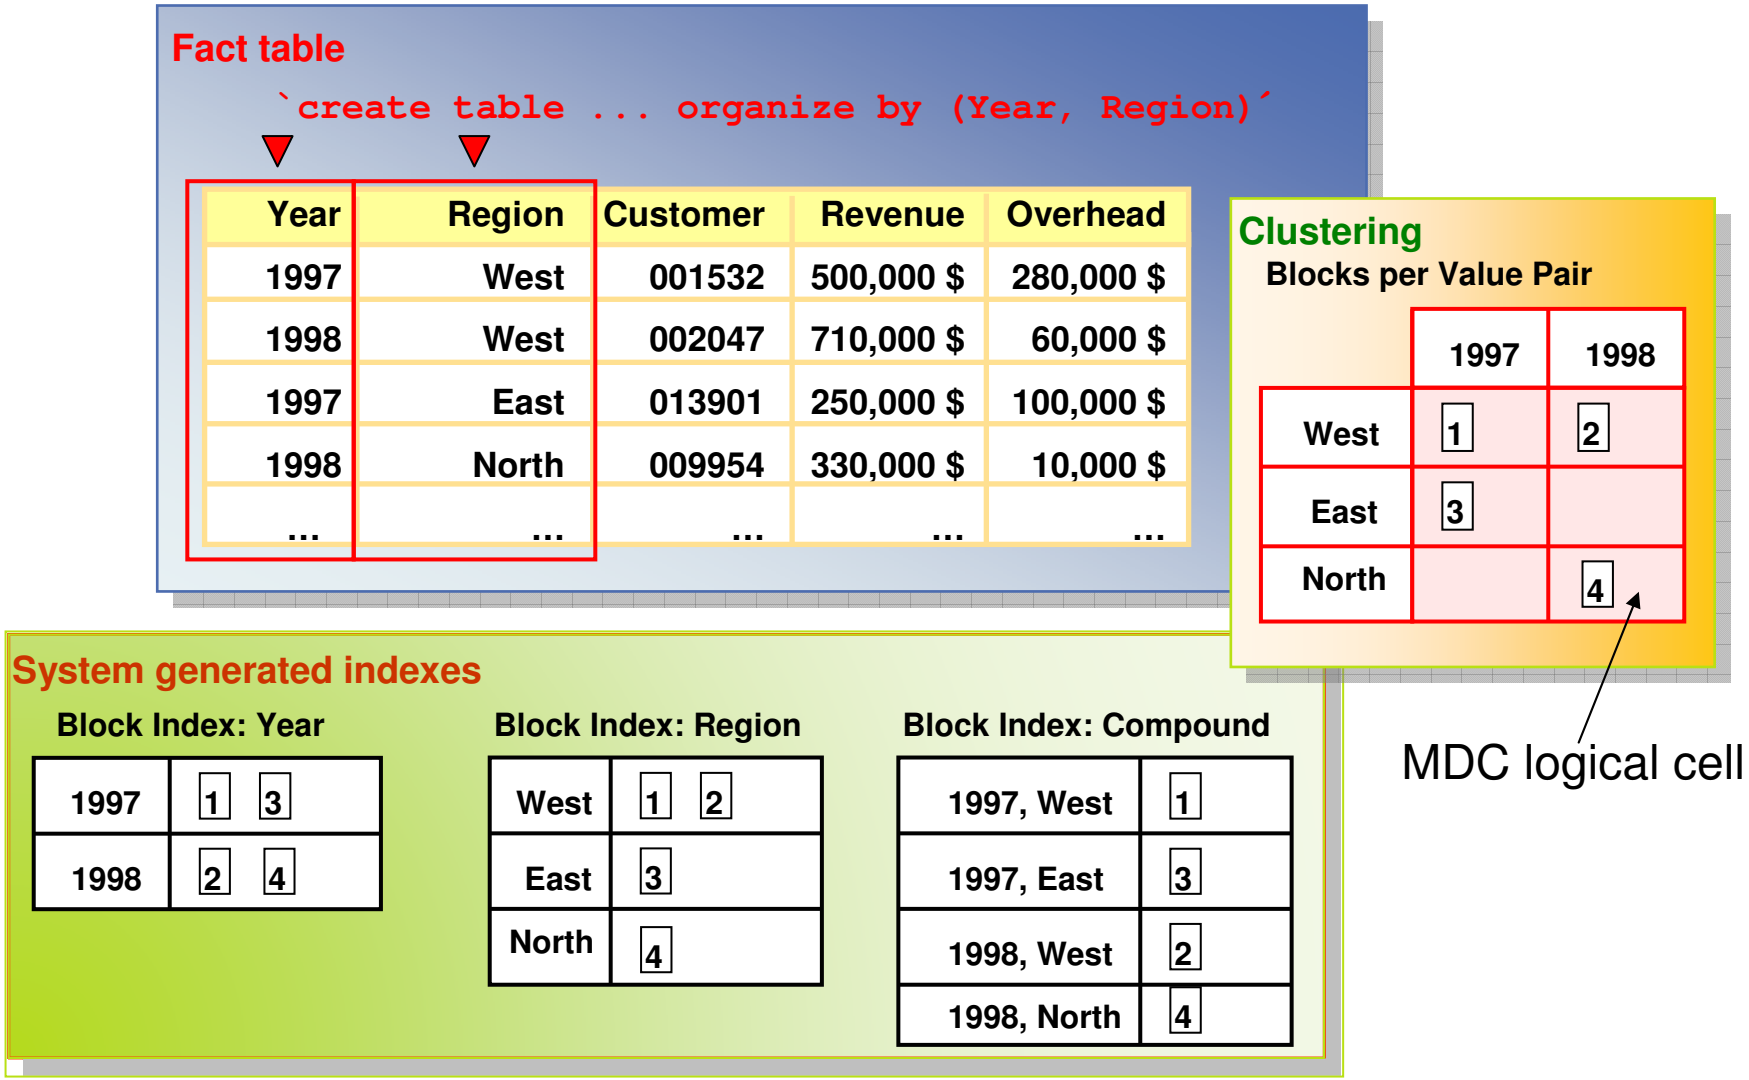

MDC Fast Insert / Fast Delete

- MDC fast insert: Reduced Locking
  - Enabled by setting LOCKSIZE parameter of an MDC table to BLOCKINSERT (during table creation or by ALTER TABLE)
  - Locks MDC blocks instead of single rows
- MDC fast delete
  - Mark pages as deleted instead of single rows
  - Cleanup of additional secondary indexes in the standard way

MDC Advantages and Considerations

- Advantages
  - Clustering is always guaranteed, no table reorganization necessary to preserve clustering order
  - Significant performance improvement for queries with restrictions on the MDC dimensions
  - Reduced maintenance overhead: block indexes are smaller and easier to maintain than row indexes
  - Support for fast data insertion and deletion (roll-in and roll-out)
- Considerations
  - Risk of high disk space consumption if wrong MDC dimensions are chosen

MDC vs. Data Partitioning

MDC	Data Partitioning
MDC cells are no separate database objects	Partitions stored in separate database objects
Query performance: MDC Cell pruning through MDC block indexes	Query performance: partition pruning
Data deletion: MDC fast delete, synchronous update of secondary indexes, space keeps allocated for table	Drop partition, asynchronous cleanup of secondary indexes, space freed up (can be used for other tables)
MDC cells added automatically	Additional partitions needed after table creation have to be added manually
Table size limits remain	Table size limits per partition

In DB2 9, data partitioning and MDC can be combined!

Supported SAP NetWeaver BI Releases

- SAP NetWeaver BI 2004s
- SAP BW 3.x, available since July 2006, with restrictions (SAP OSS note 942909)
  - Separate report RSDB6_MDC
  - SAP BW 3.0B SP 32
  - SAP BW 3.1 Content SP 26
  - SAP NetWeaver ´04 SP Stack 18 (SAP BW 3.5)
  - Requires SAP BASIS 6.20 SP 60 or SAP BASIS 6.40 SP 18

MDC Implementation Overview

- MDC is supported for PSA, DataStore objects, InfoCubes and Aggregates
- User interface for defining MDC dimensions for InfoCube Fact tables and Active Table of DataStore objects
- Re-clustering tool to convert existing InfoCubes / DataStore objects to MDC
- Disk space consumption test in SAP NetWeaver BI transaction RSRV

MDC Query Performance (internal tests)

InfoCube Details

- Fact table has 14 Million records
- 8 InfoCube dimensions
  - Time: 16 months
  - Unit: 2 units
  - Data Package: 14 packages
  - D1: 5 000 customers
  - D2: 32 700 materials
  - D3: 251 sales areas
  - D4: 2 versions
  - D5: 2 value types

Query	Restrictions
1	3 time intervals
2	3 time intervals, single customers
3	1 time interval, 1 customer interval
4	1 time interval, 1 customer interval, single materials
5	Single customer, 1 material interval
6	Single and range restrictions on all dimensions.

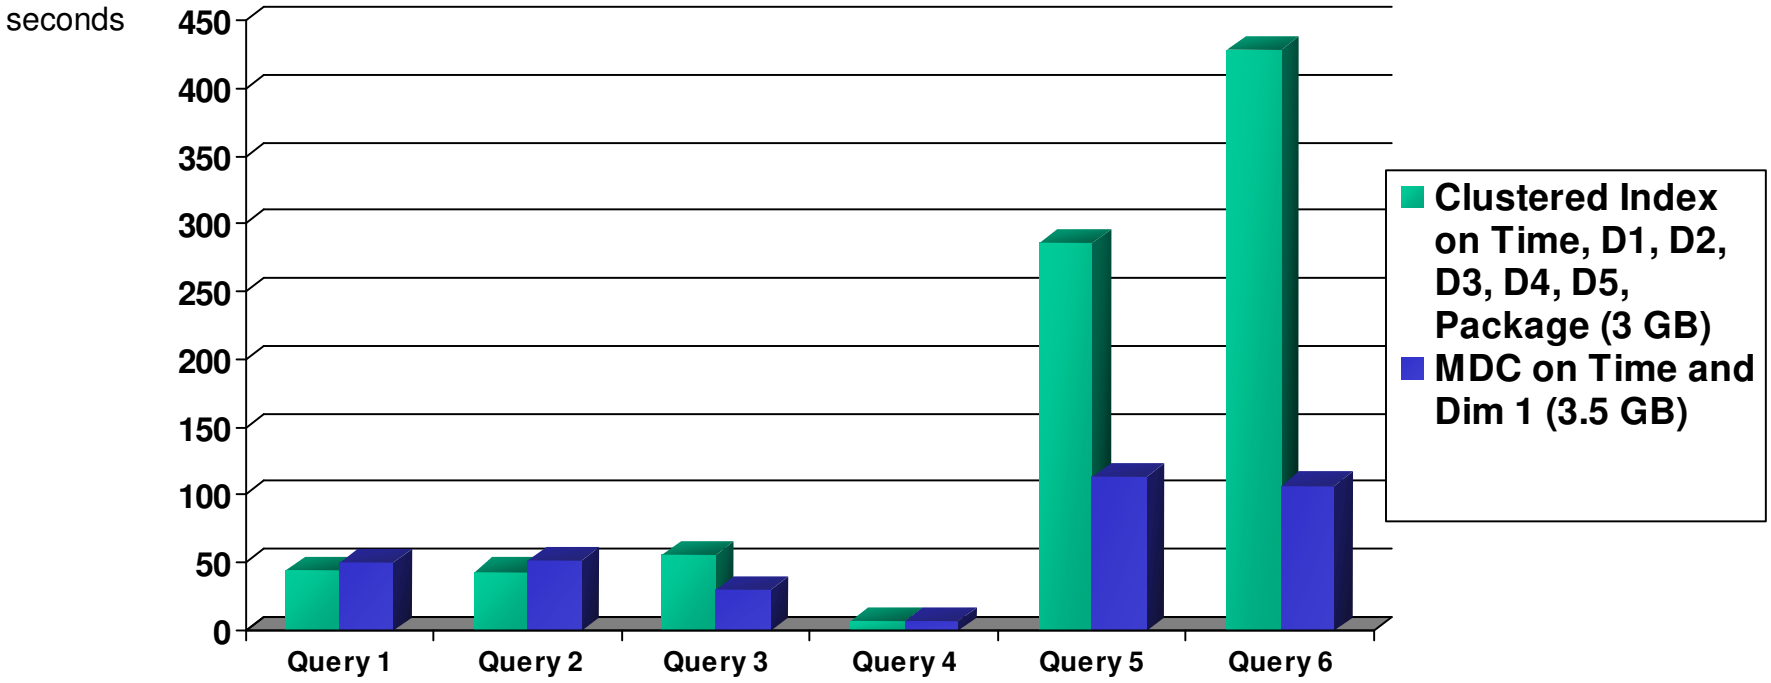

Example SAP NetWeaver BI Query (Query 6)

```
SELECT "DT"."SID_0CALMONTH" AS "S___0005" , "S1"."MATL_GROUP" AS  "K___0014" ,
        "S2"."SALESORG" AS "K___0017" , "X2"."CUSTOMER" AS  "K___0018" ,
        "D4"."SID_0VERSION" AS "S___0022" , "D5"."SID_0VTYPE"   AS "S___0023" ,
        "DP"."SID_0RECORDTP" AS "S___0002" ,  "DU"."SID_0STAT_CURR" AS "S___0009" ,
        "DP"."SID_0REQUID" AS  "S___0003" , "X2"."S__0COUNTRY" AS "S___0020" ,
        SUM (  "F"."INCORDVAL" ) AS "0INCORDVAL" , SUM ( "F"."INVCD_VAL" ) AS "0INVCD_VAL" ,
        SUM ( "F"."OPORDVALSC" ) AS "0OPORDVALSC" , COUNT( * ) AS "1ROWCOUNT"
 FROM sapr3."/MDC/FTHF_C01" "F" , sapr3."/BIC/DTHF_C01T" "DT" ,
     sapr3."/BIC/DTHF_C012" "D2" , sapr3."/BI0/XMATERIAL" "X1" , sapr3."/BI0/SMATL_GROUP" "S1" ,
     sapr3."/BIC/DTHF_C013" "D3" , sapr3."/BI0/SSALESORG" "S2" , sapr3."/BIC/DTHF_C011" "D1" ,
     sapr3."/BI0/XCUSTOMER" "X2" , sapr3."/BIC/DTHF_C014" "D4" , sapr3."/BIC/DTHF_C015" "D5" ,
     sapr3."/BIC/DTHF_C01P" "DP" ,  sapr3."/BIC/DTHF_C01U" "DU"
```

Join with Master Data and Dimension Tables

```
WHERE "F"."KEY_THF_C01T" = "DT"."DIMID" AND "F"."KEY_THF_C012" = "D2"."DIMID" AND
       "D2"."SID_0MATERIAL" = "X1"."SID" AND  "X1"."S__0MATL_GROUP" = "S1"."SID" AND
       "F"."KEY_THF_C013" =  "D3"."DIMID" AND "D3"."SID_0SALESORG" = "S2"."SID" AND
       "F"."KEY_THF_C011" = "D1"."DIMID" AND "D1"."SID_0SOLD_TO" =    "X2"."SID" AND
       "F"."KEY_THF_C014" = "D4"."DIMID" AND  "F"."KEY_THF_C015" = "D5"."DIMID" AND
       "F"."KEY_THF_C01P" = "DP"."DIMID" AND "F"."KEY_THF_C01U" = "DU"."DIMID" AND
```

Restrictions on Master Data and Dimenson Tables

```
       ((( "DT"."SID_0CALMONTH" = 200101 )) AND (( "DP"."SID_0CHNGID" = 0 ))  AND
        (( "S1"."MATL_GROUP" BETWEEN 'B00000001' AND 'B00000003' ))  AND
        (( "DP"."SID_0RECORDTP" = 0 )) AND (( "DP"."SID_0REQUID" <=  1529 )) AND
        (( "S2"."SALESORG" BETWEEN 'B003' AND 'B005' )) AND
        (( "X2"."CUSTOMER" BETWEEN 'B000002500' AND 'B000002700' )) AND
        (( "D4"."SID_0VERSION" = 1 )) AND (( "D5"."SID_0VTYPE" = 0 OR "D5"."SID_0VTYPE" = 10 ))
       )) AND "X1"."OBJVERS" = 'A' AND  "X2"."OBJVERS" = 'A'

 GROUP BY "DT"."SID_0CALMONTH" , "S1"."MATL_GROUP" , "S2"."SALESORG" ,  "X2"."CUSTOMER" ,
          "D4"."SID_0VERSION" , "D5"."SID_0VTYPE" ,  "DP"."SID_0RECORDTP" ,
          "DU"."SID_0STAT_CURR" , "DP"."SID_0REQUID" ,  "X2"."S__0COUNTRY"
```

MDC Delete Performance (I)

Test Environment
- 15 BCU units xSeries 346 (each 2-way, 8 GB memory, 6 local disks) switch connected, Linux
- DB2 9, 29 db partitions distributed on 15 BCU units, BCU node 0 has only 1 db partition

Fact table
- distributed on 14 BCU units (28 db partitions), db partartition 0 without fact data
- 33 Million rows
- different MDC configurations; row based indexes (RID) are replaced by MDC dimensions
- 2 Million rows deleted with DELETE FROM <table> WHERE <column> = <value>

  MDC tables: <column> is a MDC dimension
  Non-MDC table: clusterfactor = 1.0 on <column>

MDC Benefit
- Fewer RID index updates
- Fewer logging

Result
- MDC improves Delete Performance
- Delete Stmt runs 33% faster, if 3 RID indexes are replaced by 3 MDC block indexes

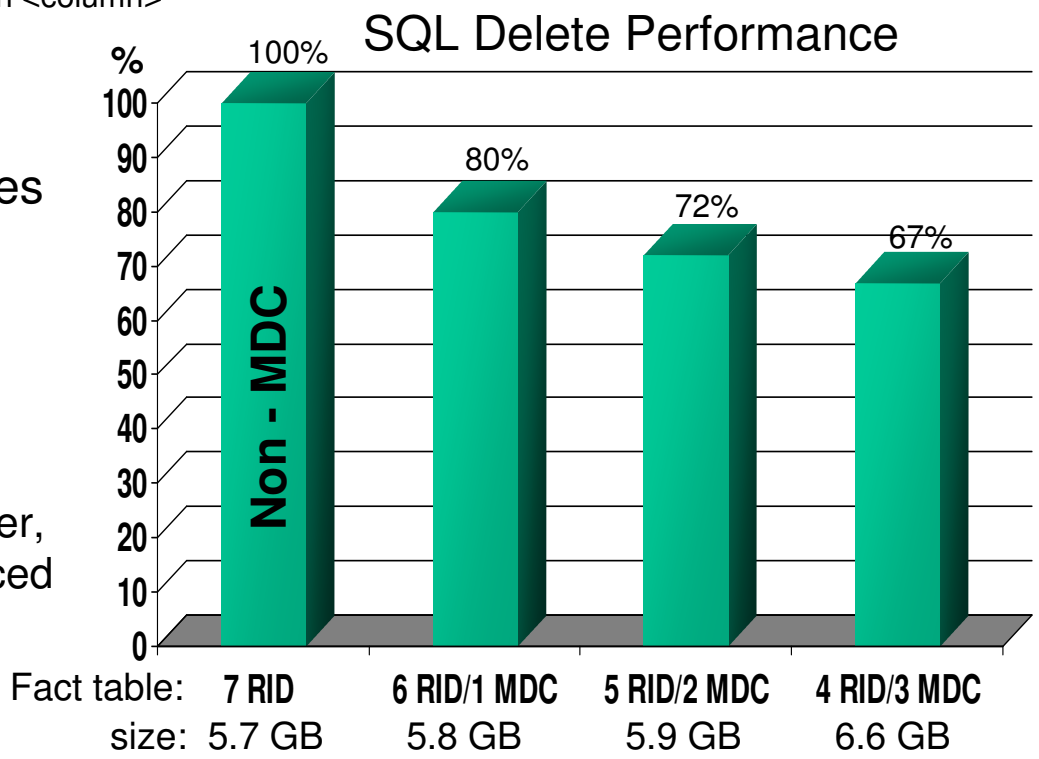

MDC Delete Performance (II)

Test Environment
- 1 BCU unit xSeries 346 (2-way, 8 GB memory, 6 local disks), Linux
- DB2 9, 1 db partition, 1.5 GB Bufferpool

Fact table
- 10 Million rows
- different MDC configurations; row based indexes (RID) are replaced by MDC dimensions
- 1 Million rows deleted with DELETE FROM <table> WHERE <column> = <value>

  MDC tables: <column> is a MDC dimension

  Non-MDC table: clusterfactor = 1.0 on <column>

MDC Benefit
- Fewer RID index updates
- Fewer logging

Result
- MDC improves Delete Performance
- Delete Stmt runs 37% faster, if 3 RID indexes are replaced by 3 MDC block indexes

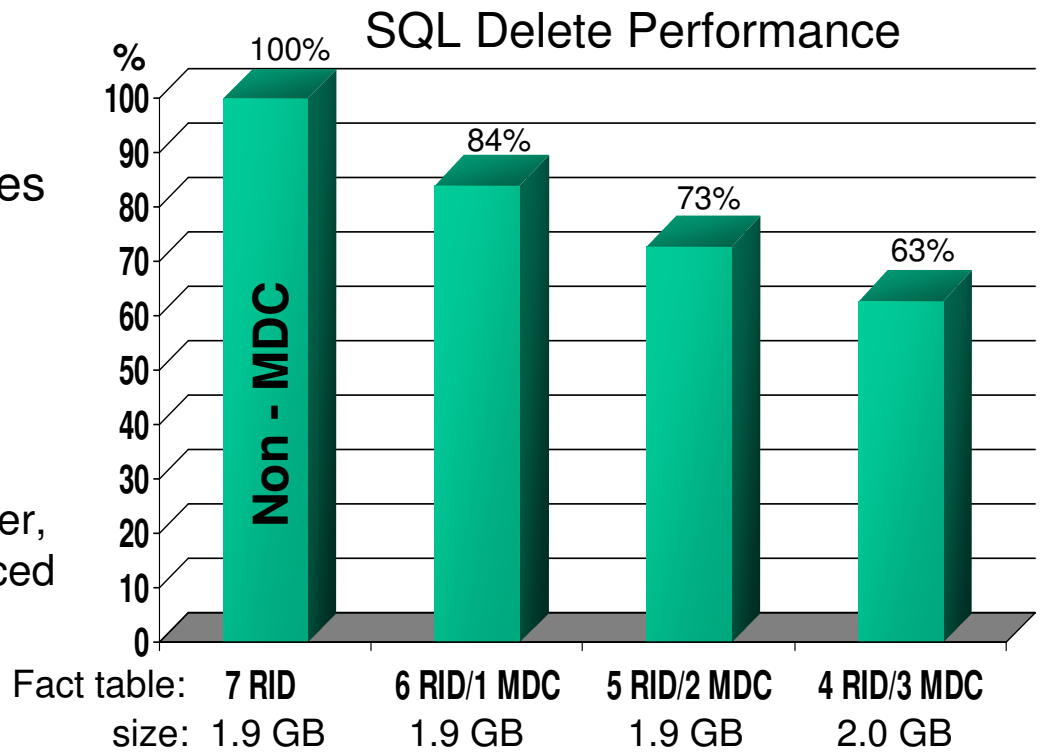

MDC Insert Performance

Test Environment
- 15 BCU units xSeries 346 (each 2-way, 8 GB memory, 6 local disks) switch connected, Linux
- DB2 9, 29 db partitions distributed on 15 BCU units, BCU node 0 has only 1 db partition

Fact tables
- distributed on 14 BCU units (28 db partitions), db partartition 0 without fact data
- 31 Million rows
- different MDC configurations; row based indexes (RID) are replaced by MDC dimensions
- 2 Million rows inserted in chunks of 100 000 records with INSERT INTO <table> VALUES (?, ?, … ?)
  MDC tables: block locking switched on

MDC Benefit
- Fewer RID index updates
- Fewer logging

Result
- MDC improves Insert Performance
- INSERT Stmts run 11% faster, if 3 RID indexes are replace by 3 MDC block indexes

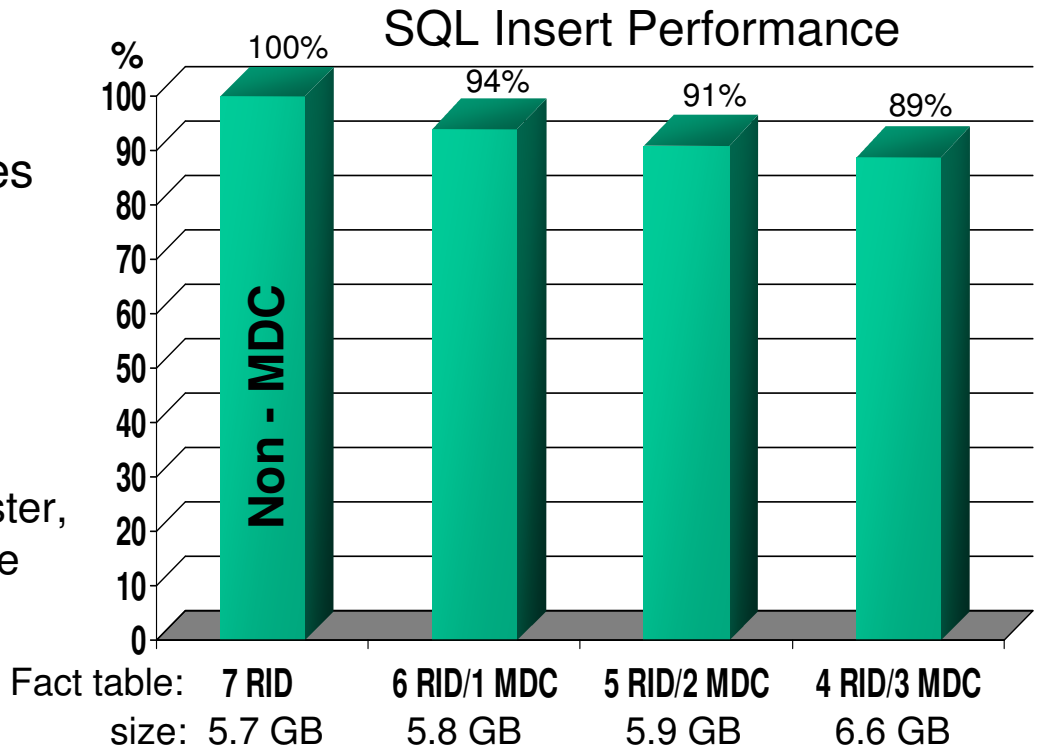

Customer POC: Environment

- SAP BW 3.5
- Database Server on AIX 5.2 pSeries Power PC Power 4 (2 way, 8 GB main memory)
- Database:
  - DB2 LUW V8.1 FP9, 64 Bit
  - Fact tables are distributed on 5 partitions
- Sales Statistics InfoCubes:
  - 15 dimensions: 3 default (time, package, unit), 12 user defined dimensions

Dimension	Cardinality	Dimension	Cardinality
DT Time (days)	667		
DTF Time (fiscal period)	36	D6 Type	198
DU Unit	45	D7 Supplier	63
DP Data Package	453	D8 Material 2	5 540
D1 Organization	9 294	D9 Customer 2	21 598
D2 Material 1	6 313	D10 Source System	4
D3 Version	25	D11 Market Place	1
D4 Customer1	26 354	D12 Ship-to Party	16 958
D5 Channel	66		

  - Fact table with 32 Million records
    - Clustered Index on DT, D1, D2, …, D12, DU, DP
  - Record size: 1 KByte
  - Fact table size: 32 GB

Customer POC: SAP BW Query Analysis

- Five most important SAP BW Queries (Sales Trend Reports) are analysed.
- Restrictions for each dimension are analysed to figure out suitable MDC dimensions
- Table below shows, how many rows are filtered by each dimension
  (T = Thousand, M = Million)

SAP BW Query	Selected Rows	DP	DT	DU	D2	D3	D4	D5
Q1	108 T	30 M	15 M	20 M	355 T	15 M	914 T	
Q2	170 T	30 M	15 M	20 M	11.5 M	15 M	914 T	
Q3	94 T	30 M	723 T	20 M	29 M	15 M	25 M	27.5 M
Q4	5.2 M	30 M	22.5 M	20 M		15 M	24 M	30 M
Q5	1.2 M	30 M	13 M			13 M	26 M	

Customer POC: Disk Space Analysis

Size of allocated disk space depends on chosen MDC dimensions

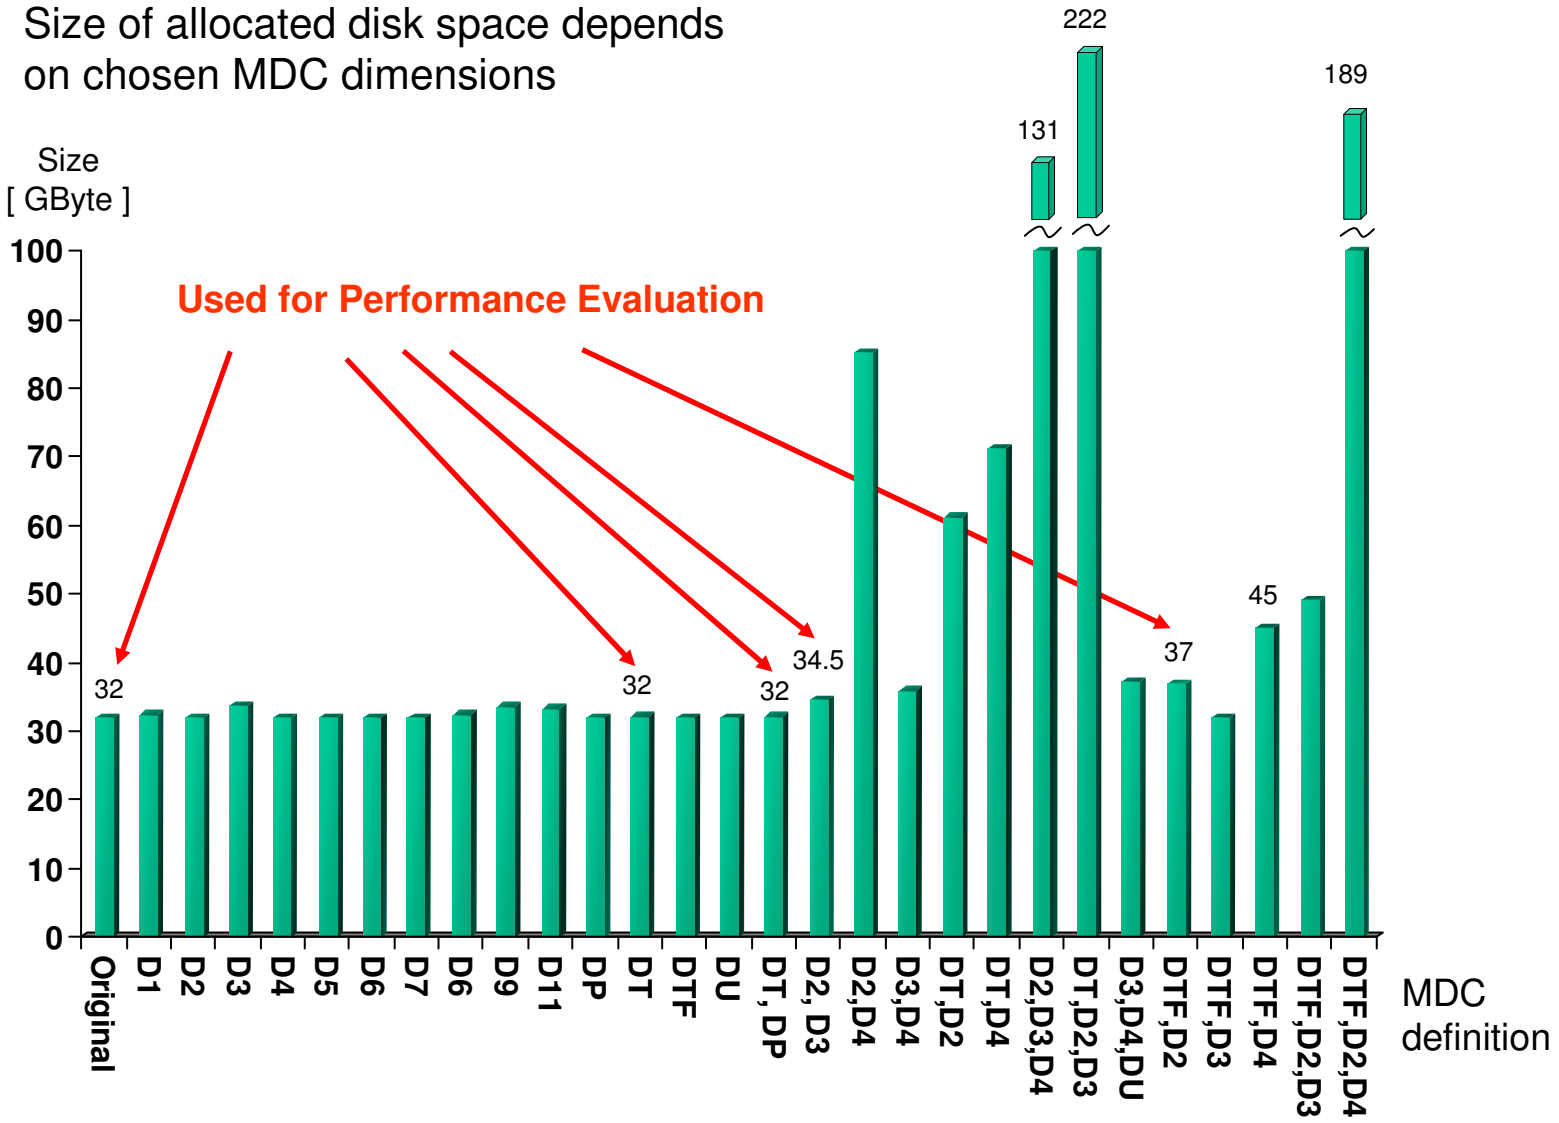

Customer POC: SAP BW Query Performance

Query Characteristics

SAP BW Query	strong restrictions on	weak restrictions on
Q1	D2, D4	DT(DTF), DU, D3
Q2	D4	DT(DTF), DU, D2, D3
Q3	DT (DTF)	DU, D2, D3, D4, D5
Q4		DT(DTF), DU, D3, D4
Q5		DT(DTF), D3, D4

Query Run Time

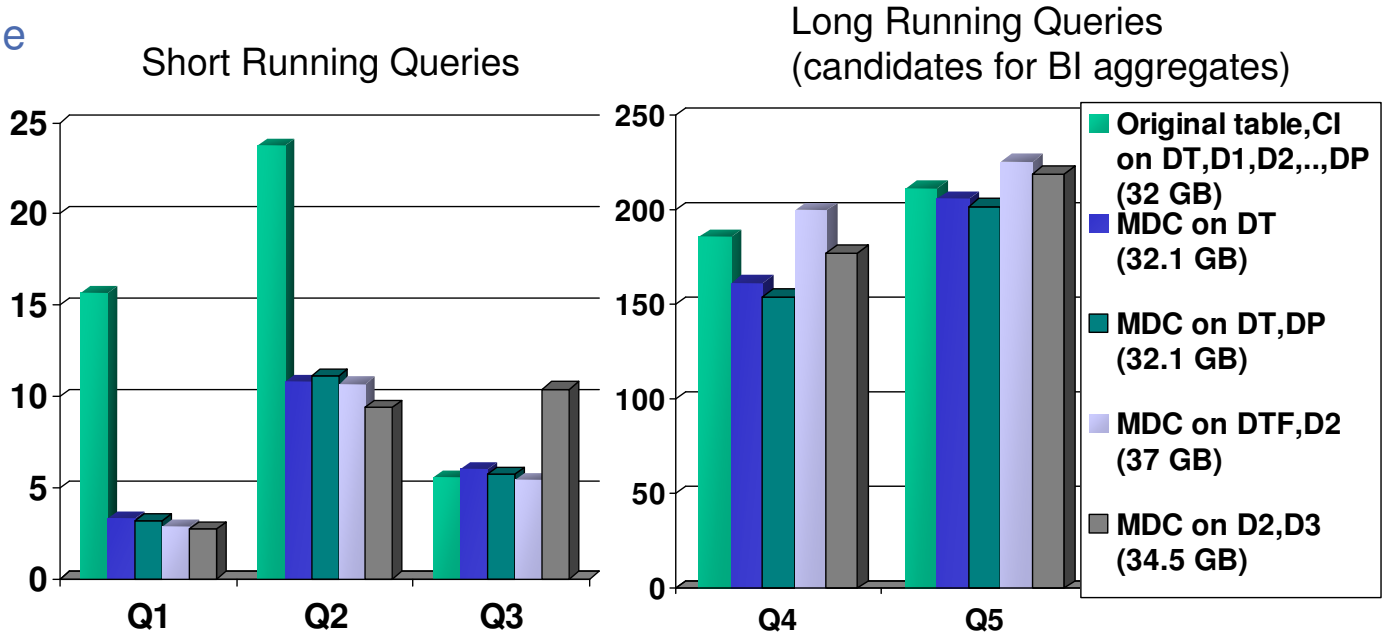

Conclusion:
- Queries run faster, if MDC table definition matches query restrictions; up to factor 5.7 faster (Q1 on MDC with D2,D3)
- One query run slower, because table definition does not match with query restrictions; factor 1.8 slower with Q3 on MDC with D2,D3
- Less run time differences for long running queries independent from MDC definition.

La compression

Row Compression Using a Compression Dictionary

- Repeating patterns within the data (and just within each row) is the key to good compression. Text data tends to compress well because of reoccurring strings as well as data with lots of repeating characters, leading or trailing blanks

Name	Dept	Salary	City	State	ZipCode
Fred	500	10000	Plano	TX	24355
John	500	20000	Plano	TX	24355

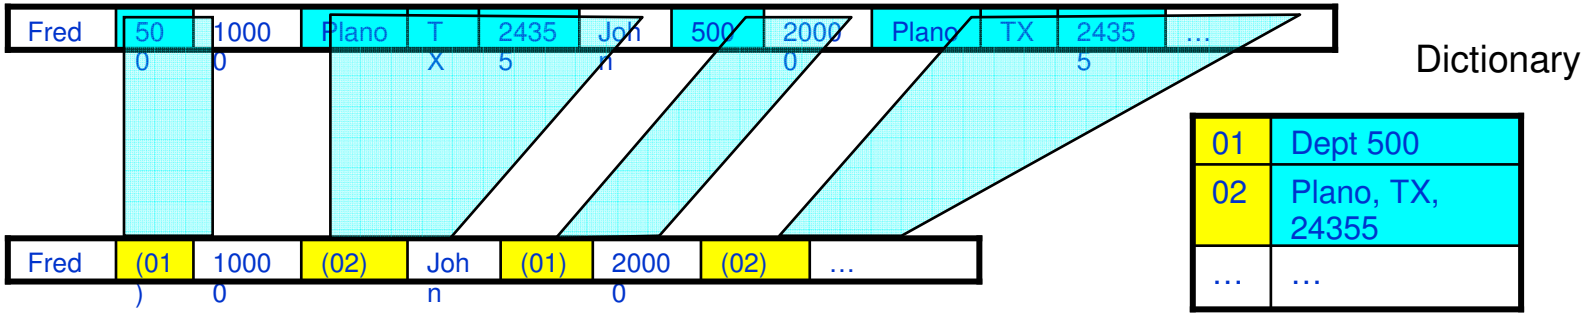

More Compression Ratios (Customer Data)

Compression Type	32KB Page Count	Space Required on Disk
No compression	5893888	179.9GB
Row compression	1392446	42.5GB

% Pages Saved: 76.4%

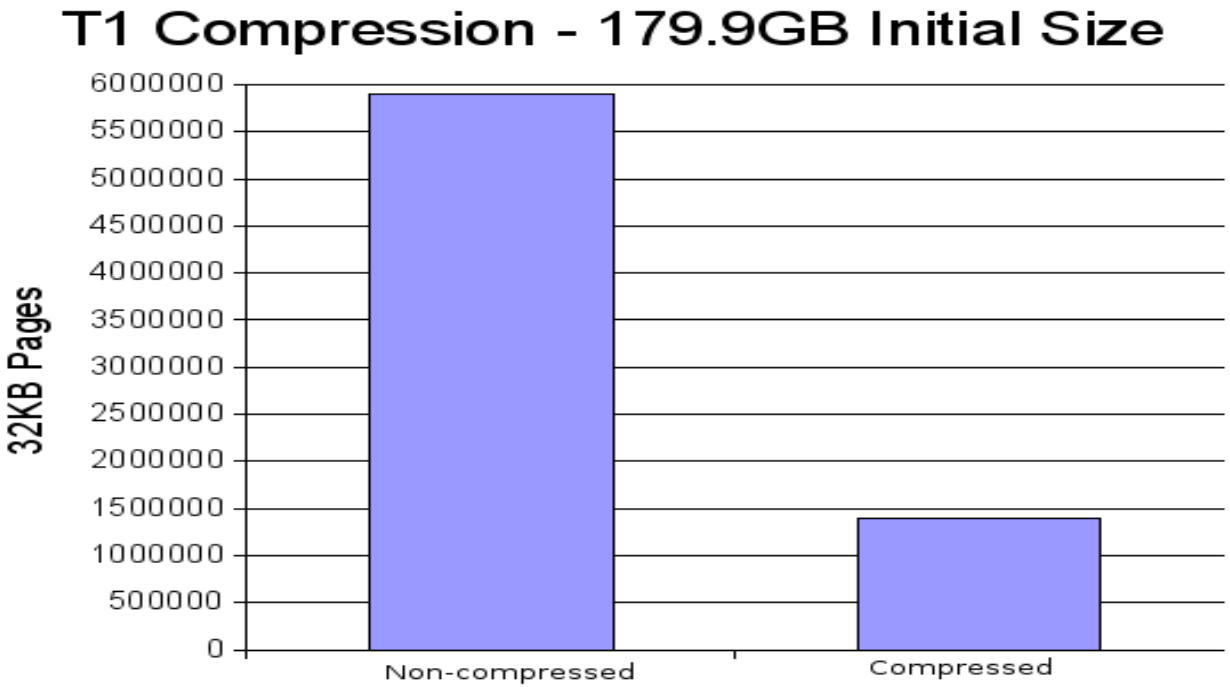

Compression Ratio – Customer Data

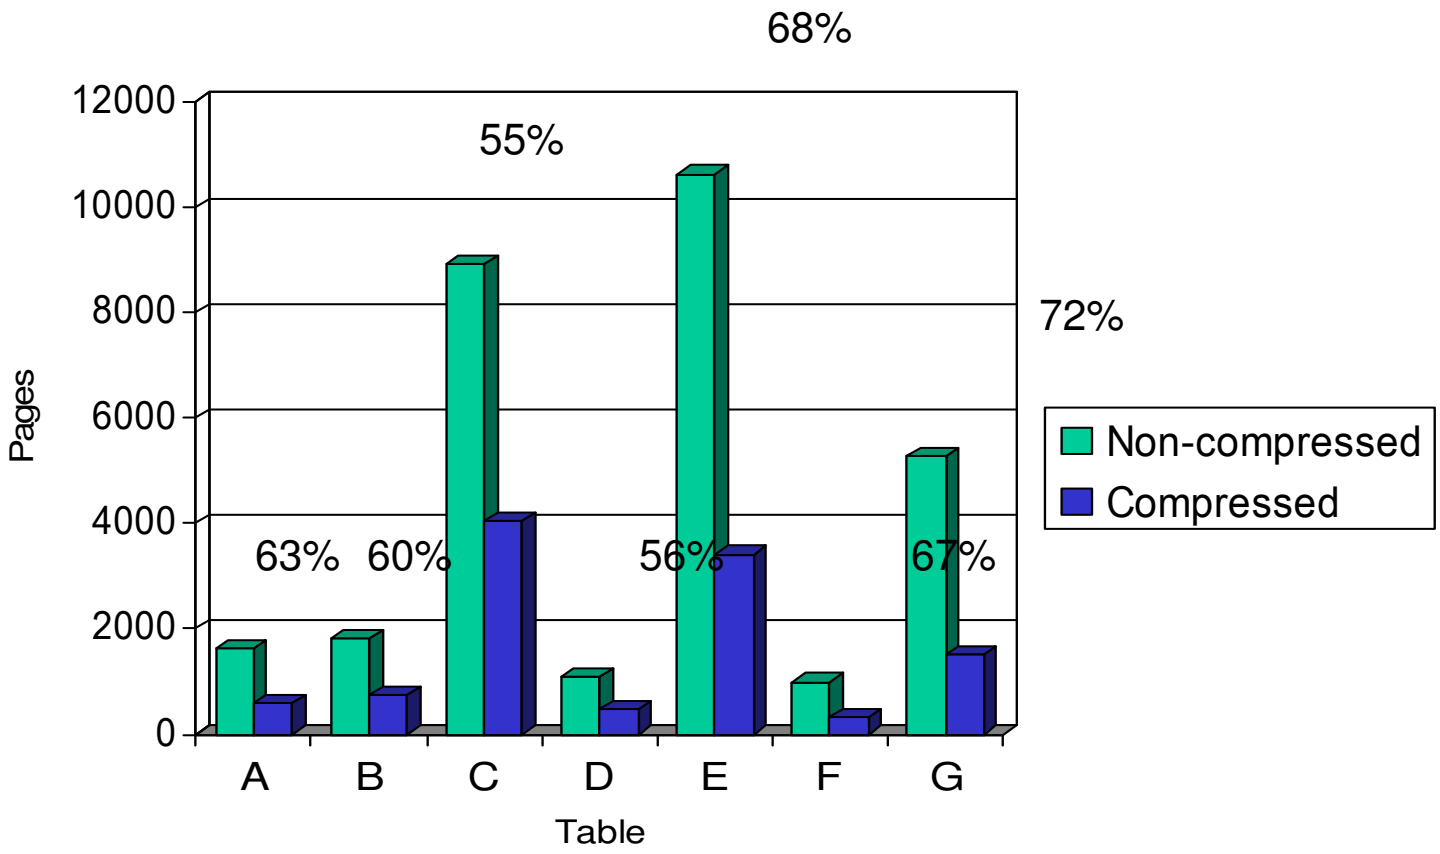

Bien plus que des économies disques

- Le dispositif de compression offre bien des avantages à tous les clients ayant des bases de données volumineuses. Le premier avantage est bien sur la possibilité de diminuer la taille des données sur disque et les besoins d'infrastructure de stockage. Mais au-dela de ce premier niveau de bénéfice, il faut y ajouter :

  - De meilleures performances car la compression réduit le nombre d'entrée sortie nécessaires pour récupérer les données.
  - Une gestion plus efficace de la mémoire car le cache des données compressées est plus efficace.
  - La diminution des images de sauvegarde de moitié et donc la diminution des temps de sauvegarde dans une même proportion.
  - La diminution des temps de restauration liée à la diminution de la taille de l'image de sauvegarde
  - Les économies d'énergie , de coût de refroidissement et de mètres carré liées à la diminution de l'infrastructure disques
  - La réduction du trafic réseau si vous utilisez HADR

De meilleures performances globales du système

La diminution du volume des données permet de diminuer les opérations d'attente sur entrées/sorties économisant des cycles CPU compensant largement le surcoût lié à l'appel de la routine de compression et améliorant grandement les performances globales du système comme en témoigne les graphiques ci-dessous

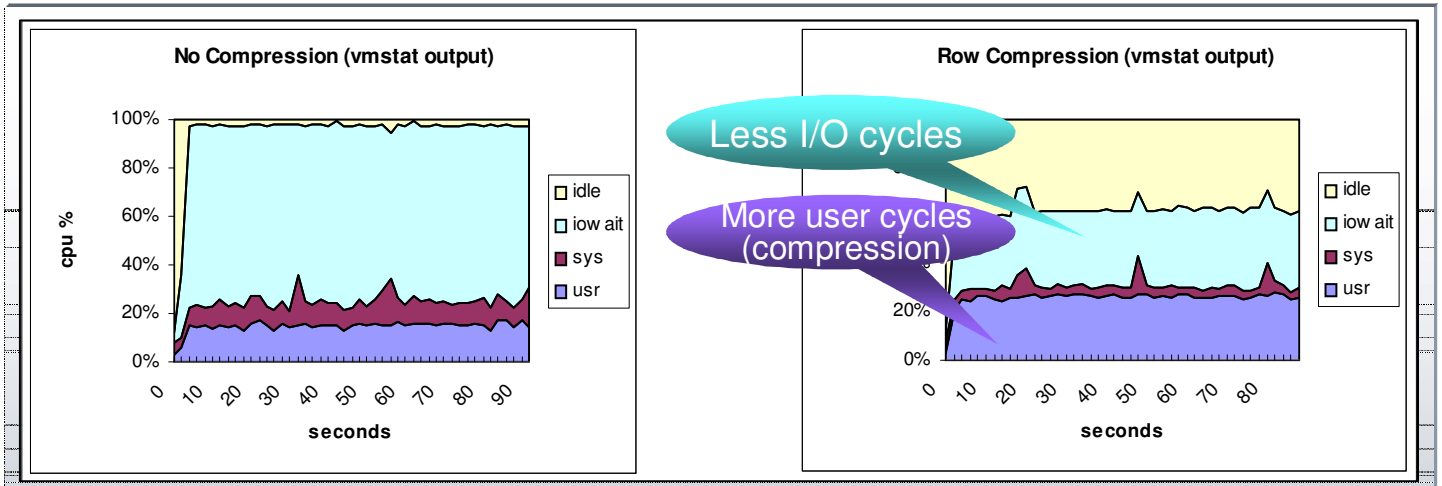

Tablescan	No Compression	Row Compression
# of Rows	59,986,052	59,986,052
# of Pages	1023162	426292
MB	7993	3330
Run time (secs)	**90.993**	**43.268**
MB/sec	87.847	76.972
Total user cpu time (sec.ms)	48.910	94.490
Time waited for prefectch	302878	67950
user cpu time/row	0.00000082	0.00000158
user cpu time/page	0.000048	0.000222
rows/page	59	141

Overall Performance Significantly Improved

Compression : diminution des temps de sauvegarde

La diminution du volume des données impacte aussi :

- **La taille des fichiers journaux**
- **La taille des images de sauvegarde , diminuant tres sensiblement les temps de sauvegarde et les temps de resatu**

Scenario	Total User Time (seconds)	Pages Used	Table Space Size (GB)	Backup Image Size (GB)
No compression	468	1510400	11.57	12
Backup Compression Only	1028	1510400	11.57	4.2
Data Row Compression Only	198	610816	4.68	4.7
Data Row and Backup Compression	662	610816	4.68	4.2

La répartition des coûts d'investissement d'un projet

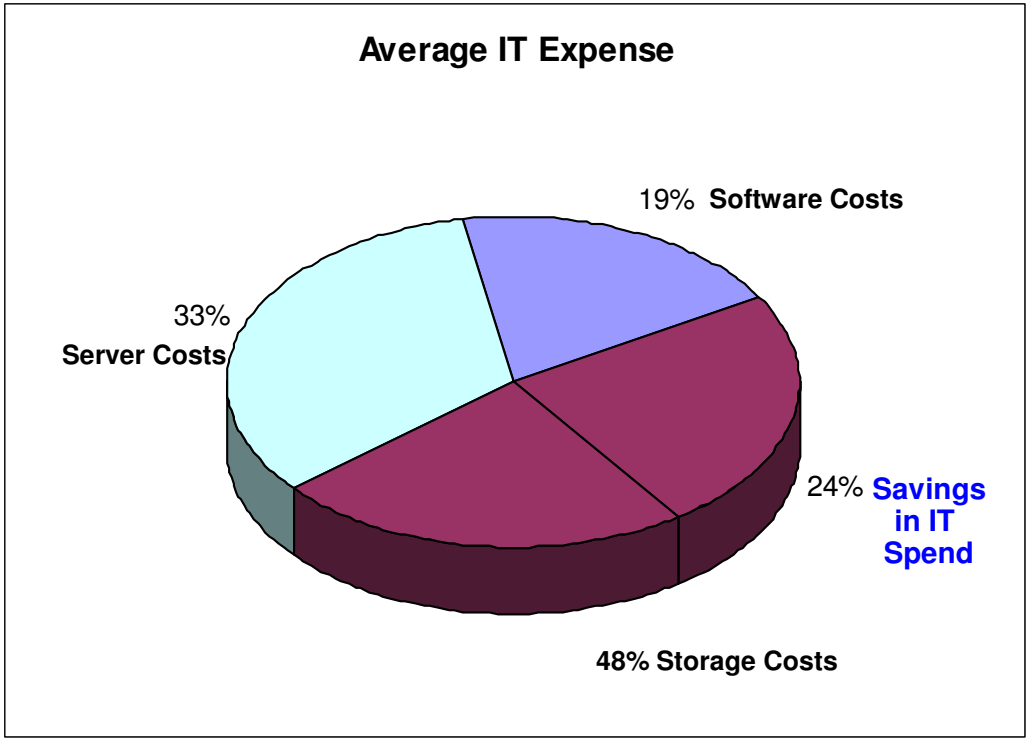

La part que représente les coûts du stockage dans un projet est conséquente , en moyenne **: 48 % du coût total ,**
Les économies apportées par la mise en sont de l'ordre de **50% ,** on peut donc économiser 24% du coût d'investissement

En résumé

SAP sur DB2: des avantages uniques

- DB2 "Database Partitioning Feature" (DPF)
  - Architecture „Shared-Nothing“seule capable d‘offrir **une scalabilté linéraire**
- DB2 "MultiDimensional Clustering" (MDC)
  - des performances des requêts SAP BI – amélioration jusqu‘à d‘un facteur 8 avec Zéro administration
- DB2 Concu pour être reconstruit
  - Capacités natives et intégrées au moteur: sauvegarde, restauration, gestion des fichiers journaux

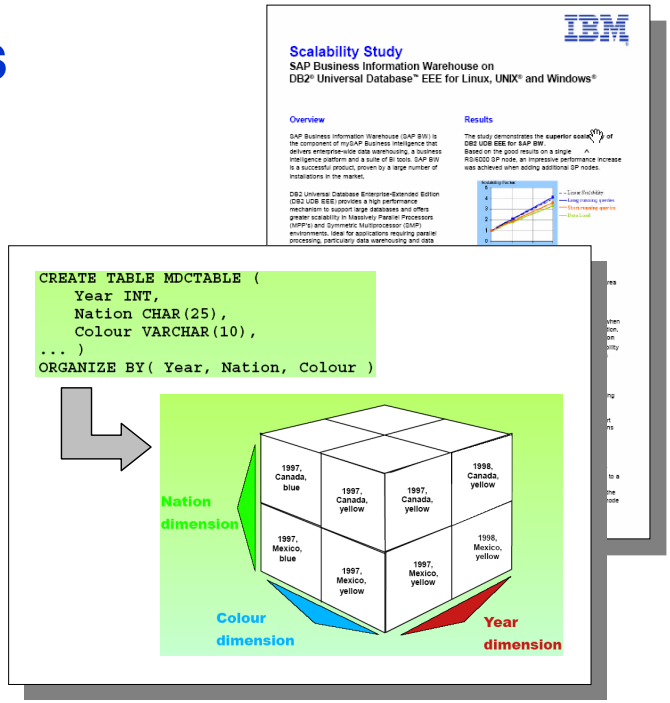

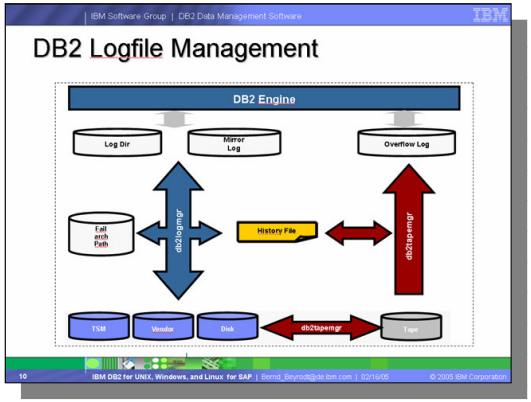

Et encore des unique DB2

- DB2 Compression
  - DB2 lead
- DB2 Load exploitation by SAP R3load
  - Temps de chargement divisé par un facteur 10
  - Complétement suporté depuis le R3load 4.6 et + (SAP Note 454173)
  - Diminution signifiactive des temps d'arrêt lors de migration
  - Available with DB2 only
- DBACockpit pour DB2
  - Performances,space, configuration , compression
  - MDC
  - Backup/Restore,Diagnostics,Alerts
  - Jobs

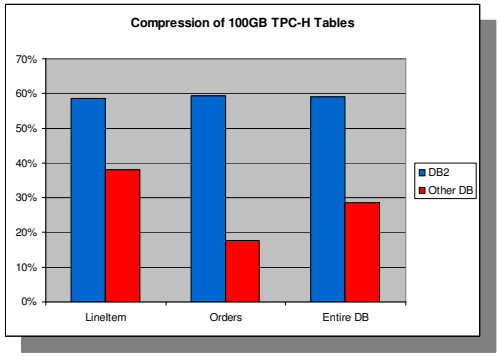

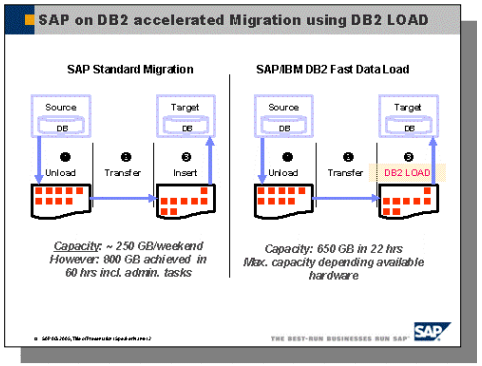

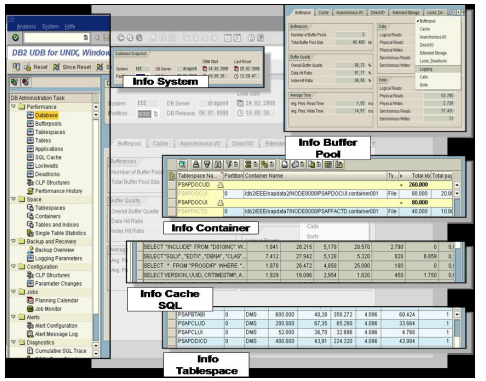

Encore un dernier dispositif

BI Accelerator – General

- High Performance Analytics (HPA) is an optional new SAP BI
- Performance Feature

- Available with SAP Netweaver 2004s

- Why BI HPA ?
  - -> Replacement of SAP BW Aggregates:
  - - Run complex queries in a few seconds
  - without BW Aggregates
  - - Reduce BW Aggregate maintenance effort
  - -> BI HPA is DB independent

Basis InfoCube

Product	Country	Customer	revenue
Desktop	Germany	Customer A	75
Notebook	Germany	Customer A	143
Desktop	Austria	Customer B	40
Notebook	Canada	Customer C	96
Printer	USA	Customer C	37
Notebook	China	Customer C	34

Aggregates

Product *

Product	revenue
Desktop	115
Notebook	273
Printer	37

Product F Notebook
Customer *

Product	Customer	revenue
Notebook	Customer A	143
Notebook	Customer C	130

SAP BI Solution without HPA

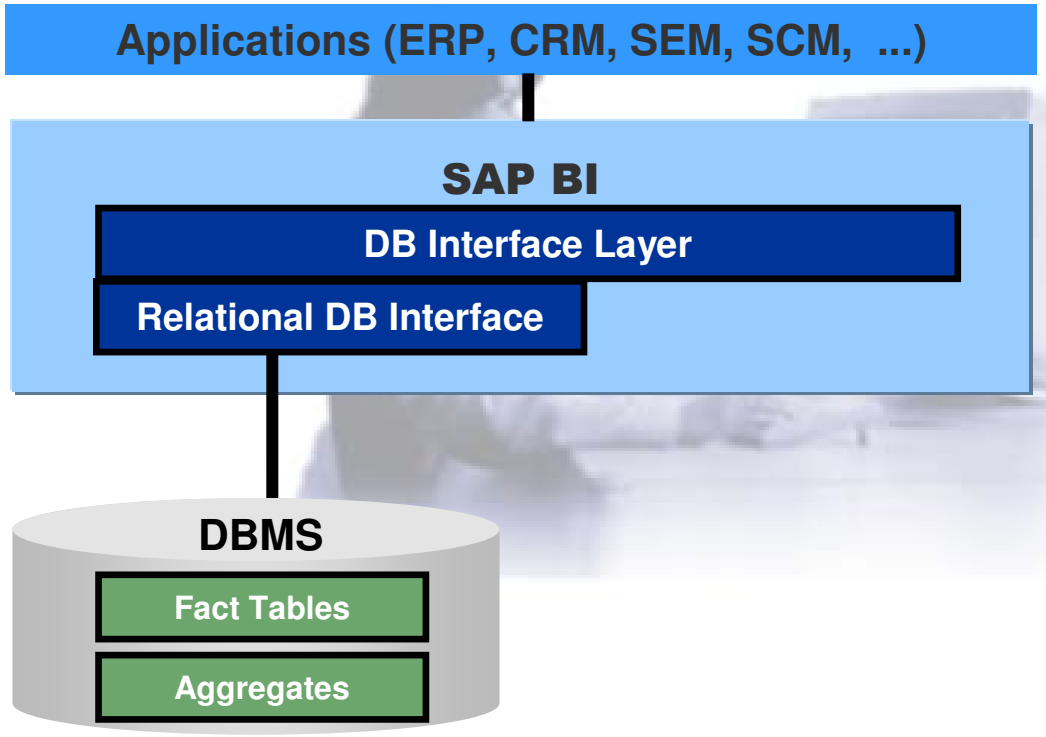

Storage + Query Processing

BI Accelerator Appliance

Increase scalability and flexibility, w/o impacting architecture

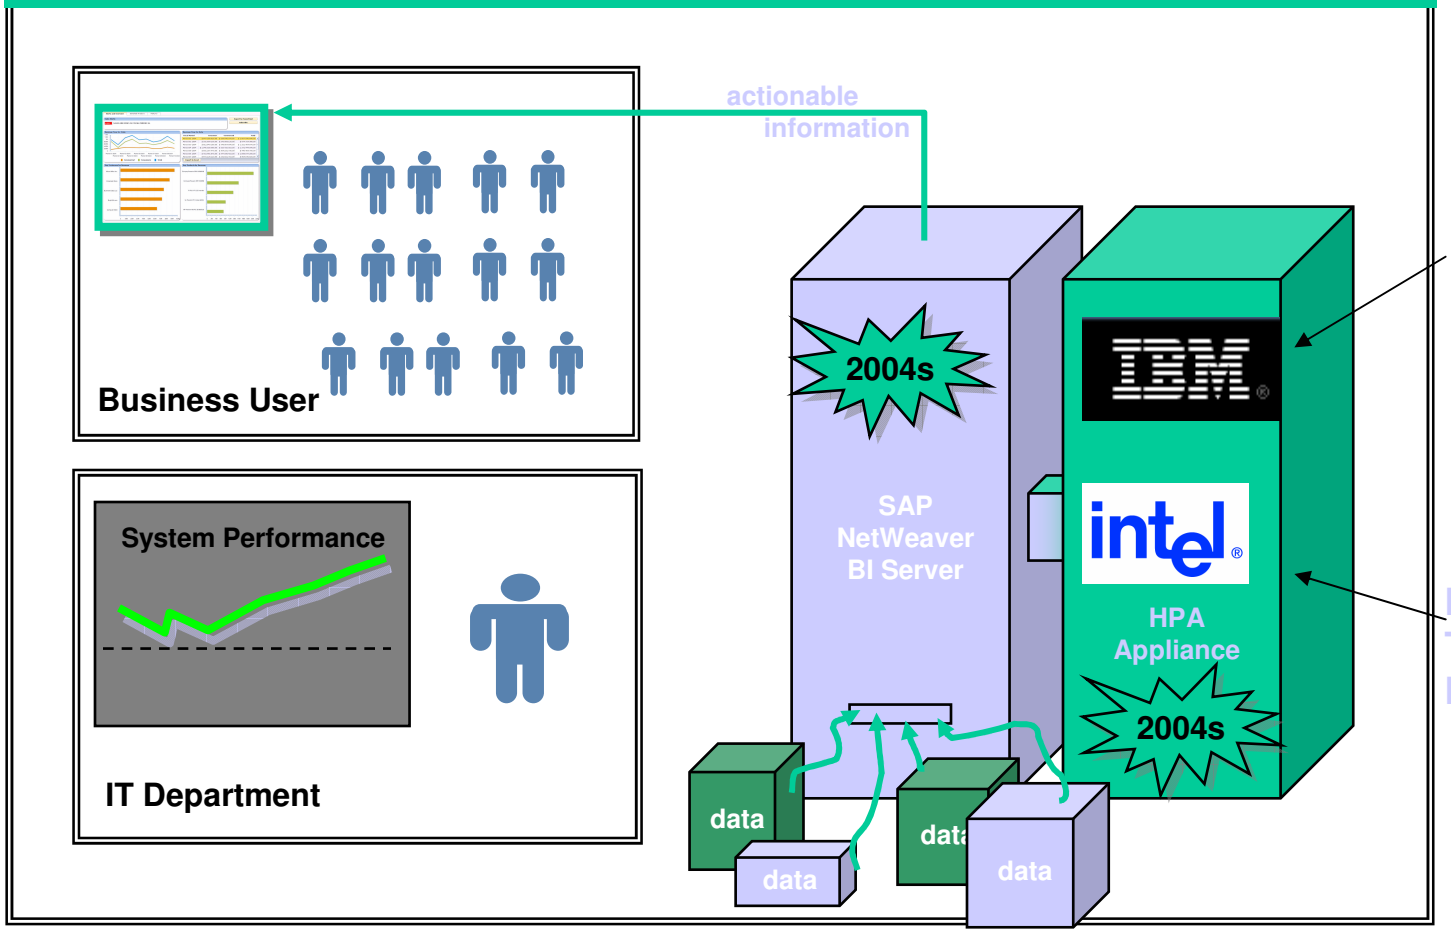

Quantum Leap in Performance plus Quantum Leap in Flexibility
= **Increased Productivity and Innovation**