Performance Optimization with SAP on DB2 for Linux, UNIX, and Windows

System Copy

Andreas Christian

IBM SAP DB2 Center of Excellence

Agenda

- Overview Heterogeneous SAP System Copy
- Standard Optimizations
- Optimizations for Migrating to DB2 LUW
- Summary

SAP System Copy - Overview

- Standard SAP System Copy
- Minimized Downtime Service (MDS)
- UNICODE Conversion

System Copy Overview – Backup / Restore

Starting with DB2 UDB Version 8, you can use redirected restore to create a Heterogeneous System Copy between two systems running on different platforms

Note 628156 - DB6: Cross-Platform System Copy using Backup/Restore with V8

Document:

"Copying Your SAP/R3 System Across Platforms Using DB2 Universal Database V8 Redirected Restore"
http://www-106.ibm.com/developerworks/db2/library/techarticle/0308nesiem/0308nesiem.html

Link:

DB2 Documentation: Backup and restore operations between different operating systems and hardware platforms
http://publib.boulder.ibm.com/infocenter/db2luw/v9r5/index.jsp?topic=/com.ibm.db2.luw.admin.ha.doc/doc/c0005960.html

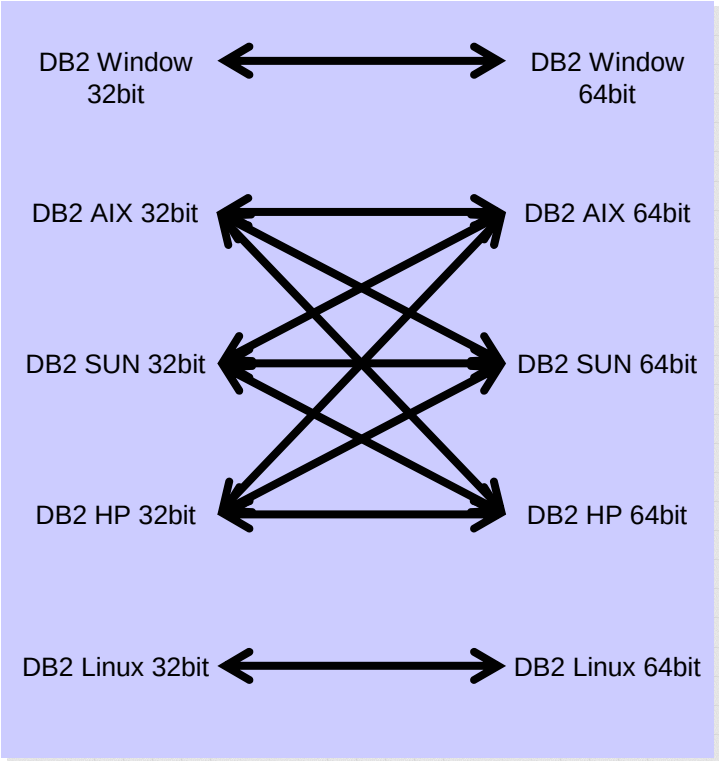

SAP System Copy Overview – R3load

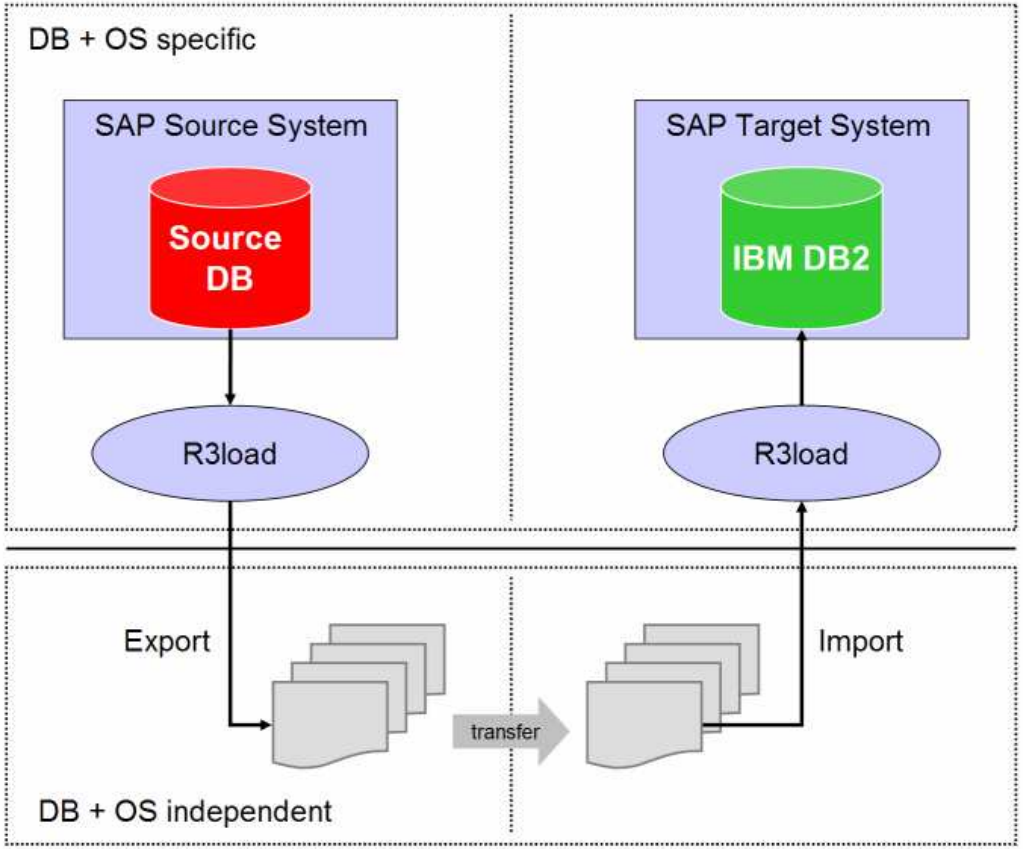

R3load features

- DB + OS independent export format
- Multiple instances of R3load can run in parallel
- DB2 Compression during Import

SAP System Copy Overview - MDS

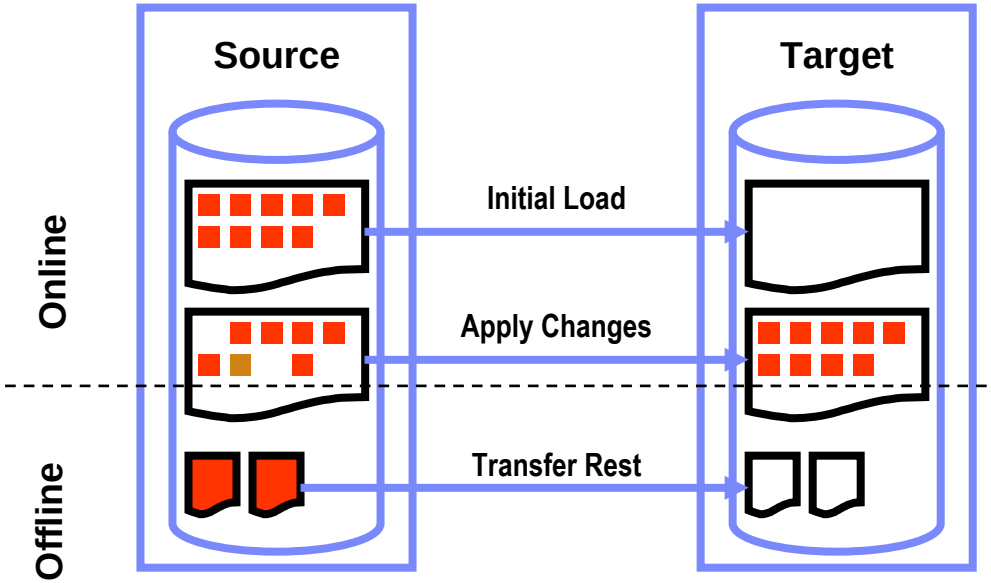

Minimized Downtime Service

- Former known as IMIG = Incremental Migration
- Fee based SAP service offering
- Currently not supported for SAP Business Warehouse
- MDS is not covered in this SAP Info Session

Required Steps

- Dirty export of large tables
- Subsequent changes on large tables are tracked via database triggers and applied to target database
- All other tables are exported/imported in the final offline phase of the MDS

UNICODE Conversion Overview

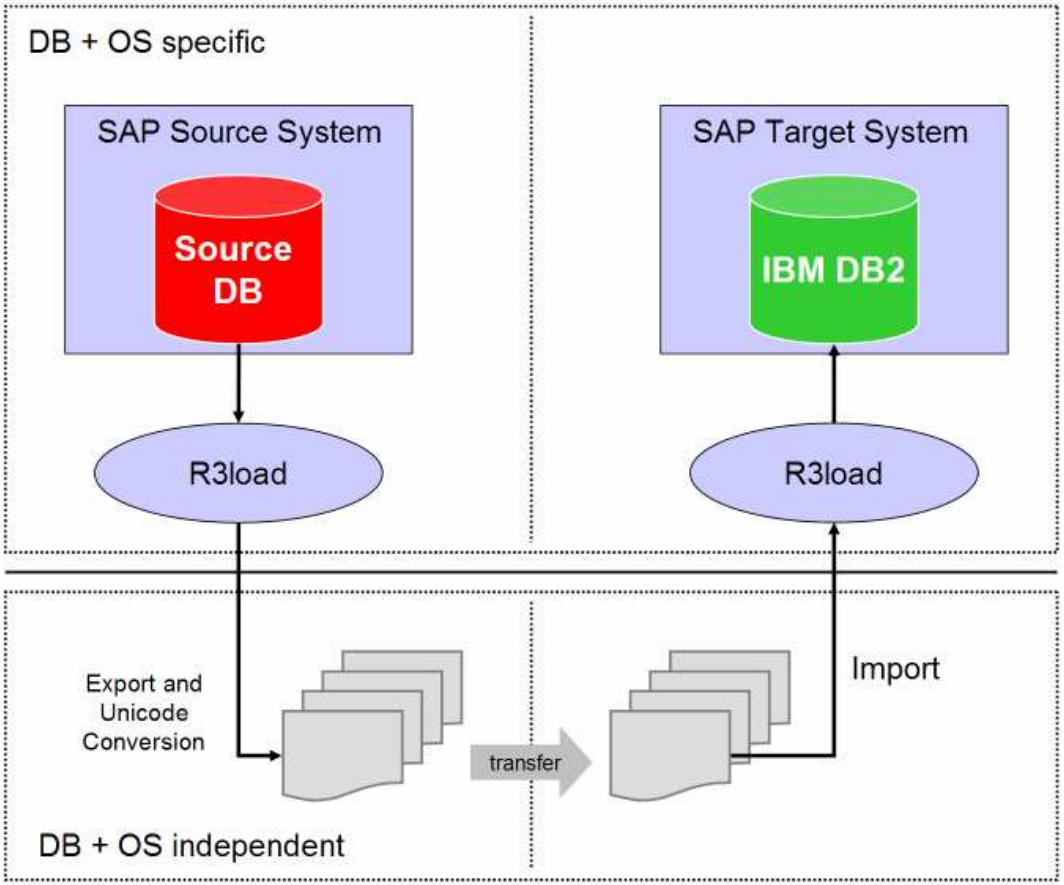

- **Conversion to UCS2 (UTF16) during export (via R3load)**
- **Conversion to UTF8 during import (by R3load or DB2 LOAD)**
- **Single code page**
  - Easy conversion
- **MDMP**
  - Complex conversion
  - Tables with text in different languages and without language key need special (manual) treatment to figure out the corresponding codepage

Standard Optimizations for SAP System Copy

- Package Splitting
- Table Splitting (usage and impact on export)
- Package Order
- Unsorted vs. Sorted Export
- Local or remote R3load
- Considerations for UNICODE Conversions
- Socket Option

Standard Optimizations for SAP System Copy

- Split packages (options top, tableLimit, packageLimit, tableFile)

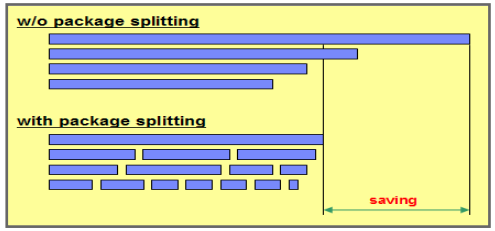

- Split large tables (important for UNICODE conversion of large table clusters)

```
R3ta -f CDCLS.STR -table CDCLS%8
```

- Start export and import of largest tables / packages first (MigMon with orderBy file)

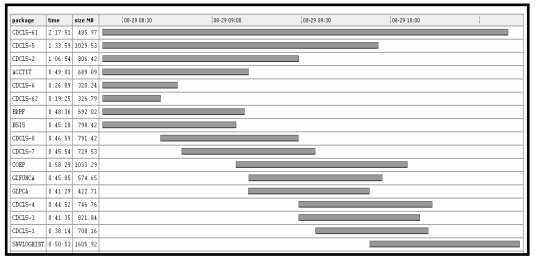

- If possible use unsorted export (remove ORDER_BY_PKEY keyword from TPL-file)

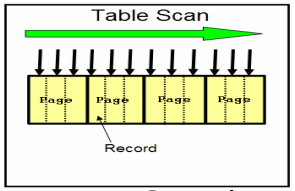

Un-Sorted

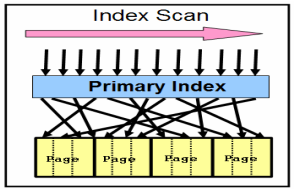

Sorted

- Use remote R3load for export if CPU is a bottleneck

- Use socket option (requires stable network connection)

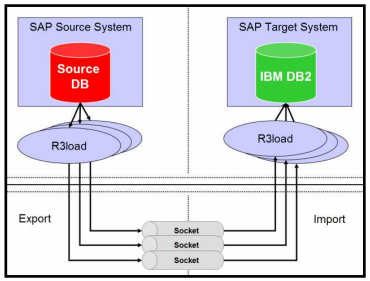

UNICODE Conversion of Table Clusters

- A **Table Cluster** is a table that stores multiple logical tables (Cluster Tables) in compressed format.
- A row of a Cluster Table can be stored in one ore more consecutive rows of the corresponding Table Cluster.

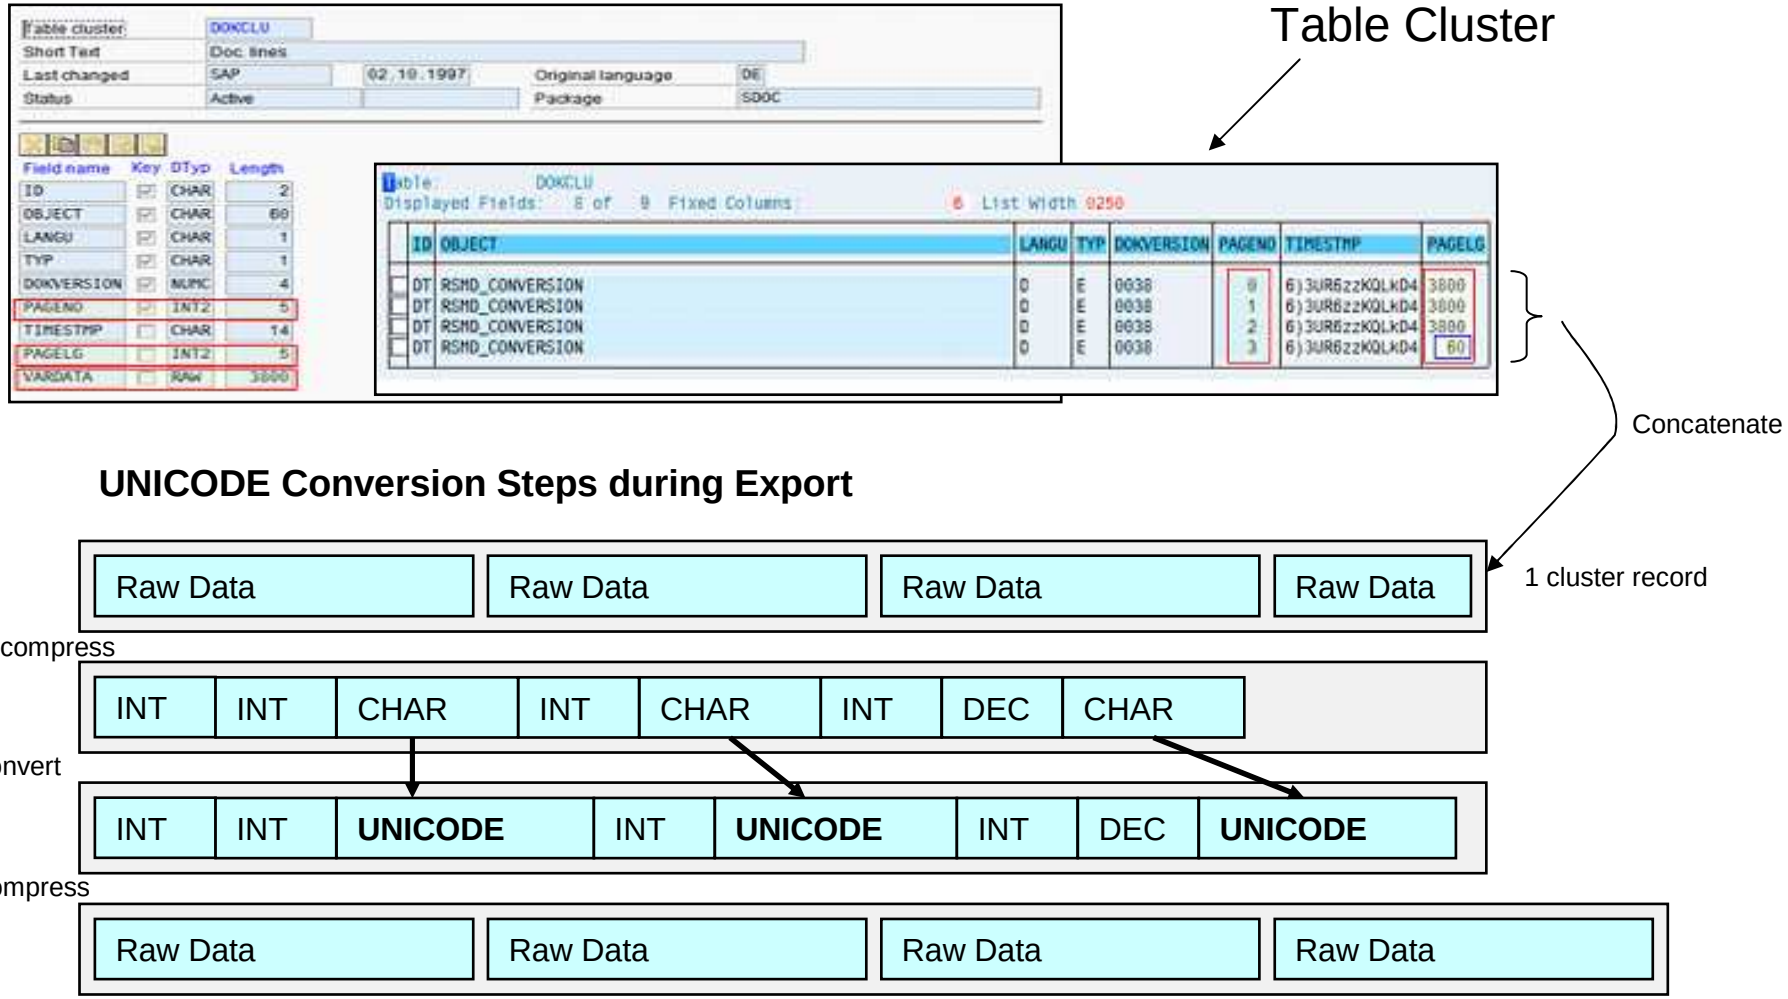

Unsorted vs. Sorted Export – Runtime Comparison

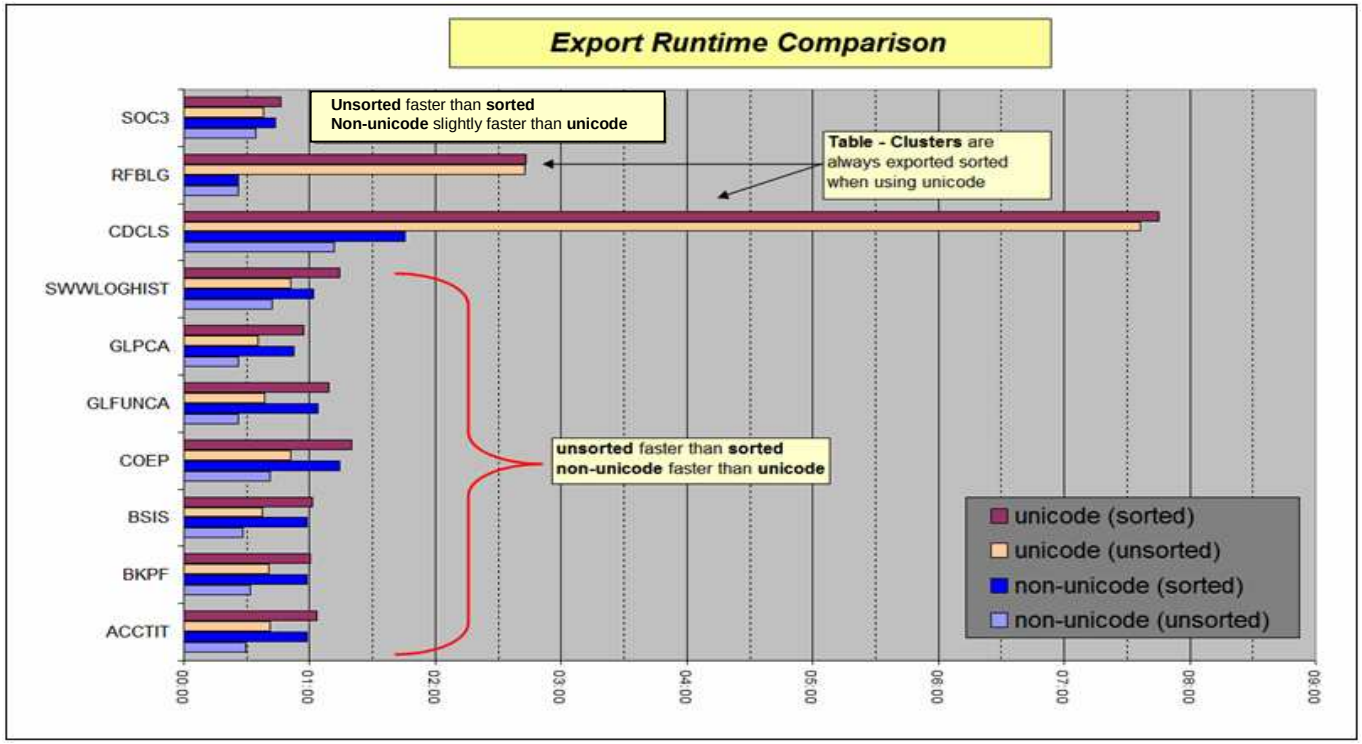

- **Unsorted export is faster than sorted export!**
- Not all table can be exported sorted. Exceptions described in SAP note 954268.
- Tables that have been exported unsorted will also be imported unsorted. This may result in insufficient access to disk on the target system (bad cluster ratio).

  => Reorg may be required on target system.

Local vs. Remote R3load

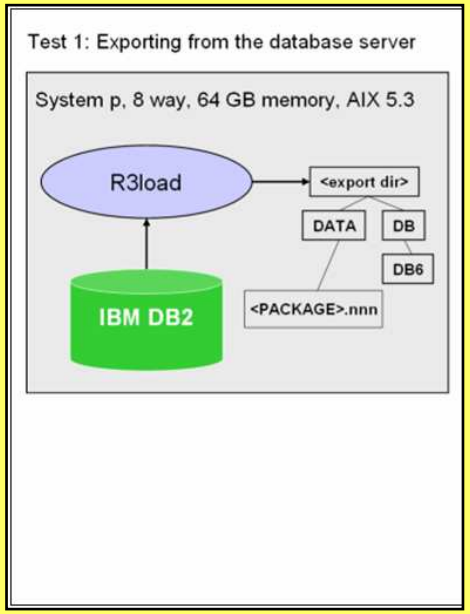

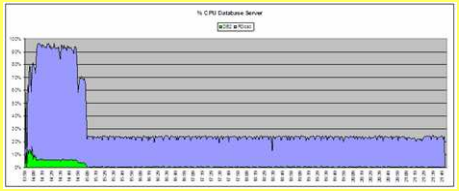

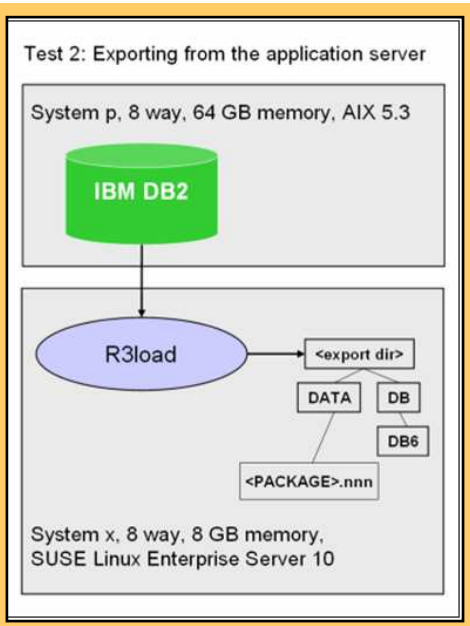

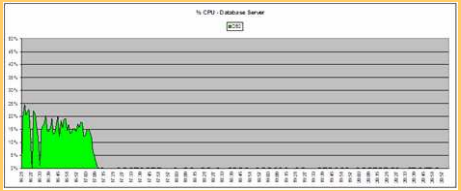

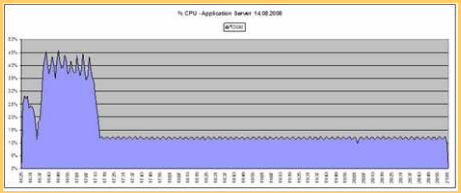

- If **CPU** is a **bottleneck** you may want to **move workload** away from the DB-server to an application server.
- Provide a **fast and stable network** connection

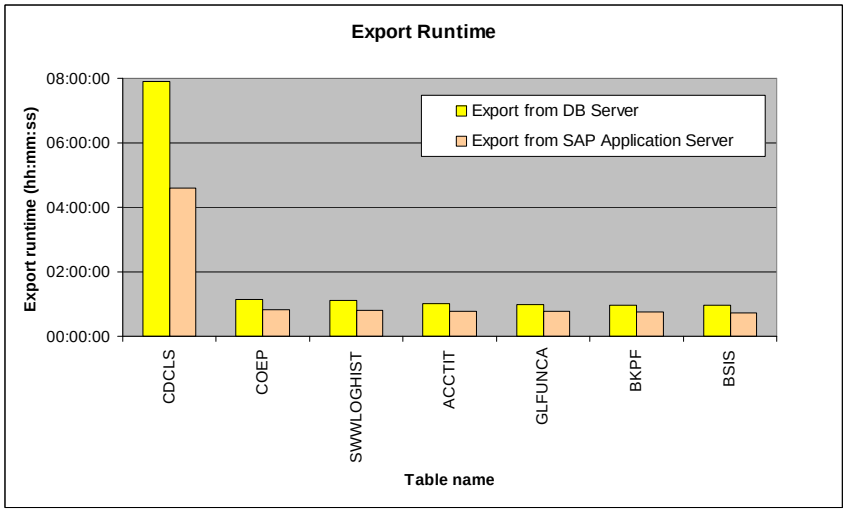

Testcase with UNICODE

Optimizations for Migrating to DB2 LUW

- DB2 Table Space Layout
- DB2 LOAD vs. INSERT
- DB2 LOAD configuration
- DB2 Memory Configuration
- Order of Index Creation
- Deployment of DB2 Compression
- Optimized DB2 Statistics

Recommended DB2 Table Space Layout

Mapping of table spaces to disks for Storage Area Networks (SAN):

- Map each LUN (logical disk) to a dedicated RAID Array
- Create one file system on each LUN

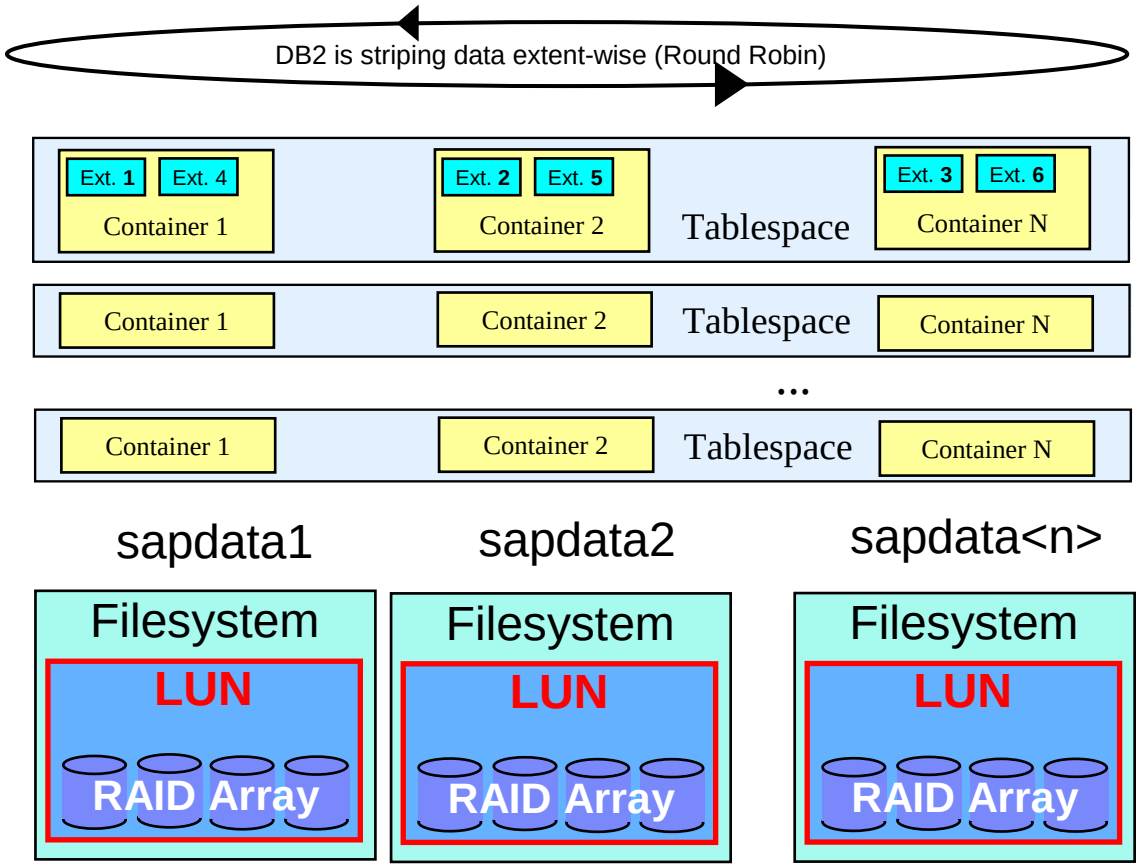

Performance recommendations:

- **Spread containers of each tablespace over all filesystems / spindels.**
- **Put each container of a table space on a separate file system**
- **Enable DB2_PARALLEL_IO for table spaces with one container per RAID device or stripe set**
- **Use 15 – 20 dedicated spindels per CPU core.**
- **Avoid too many levels of striping:**
  - DB2 stripes across containers
  - Storage controllers provide RAID striping

  => OS level striping, e.g. LVM striping should NOT be used.
- **Keep table space sizes balanced for faster backup and restore**

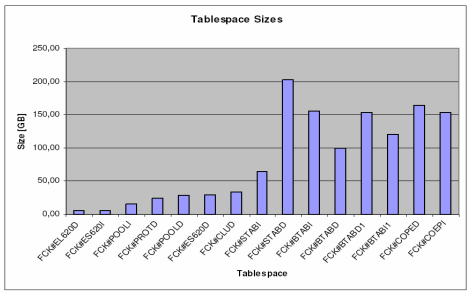

DB2 LOAD vs. INSERT - Process Model

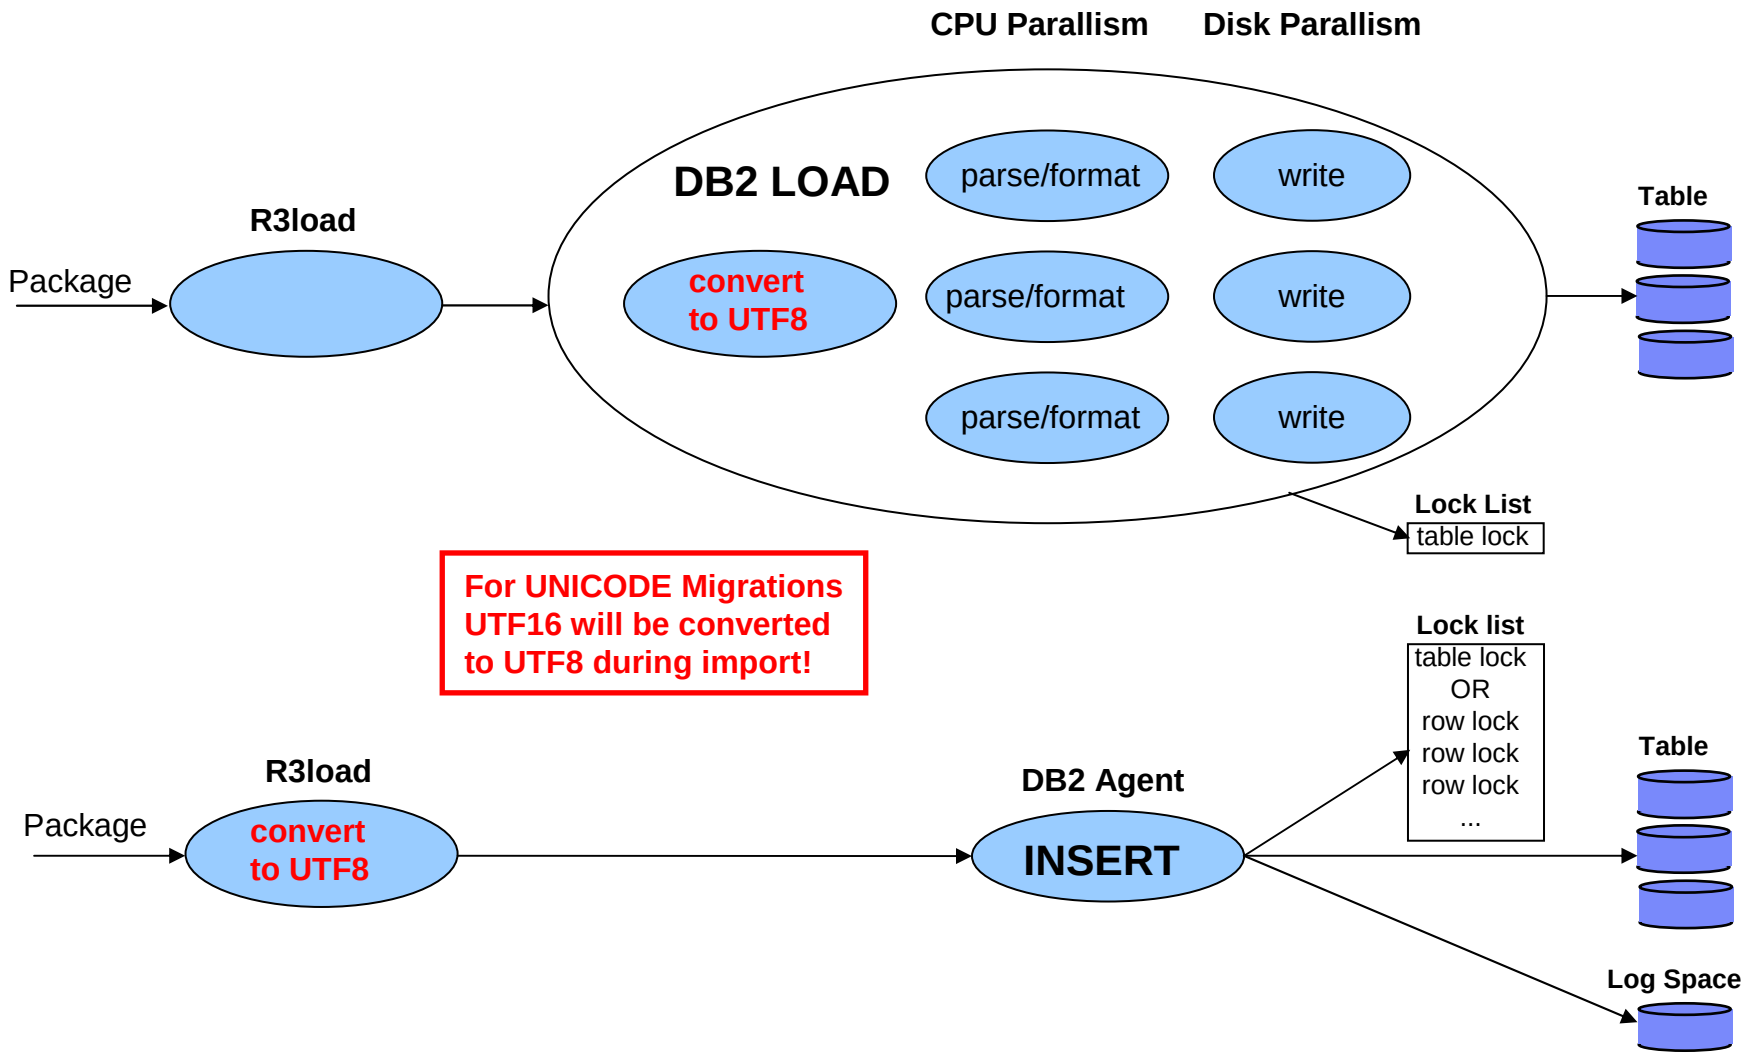

DB2 LOAD vs. DB2 INSERT

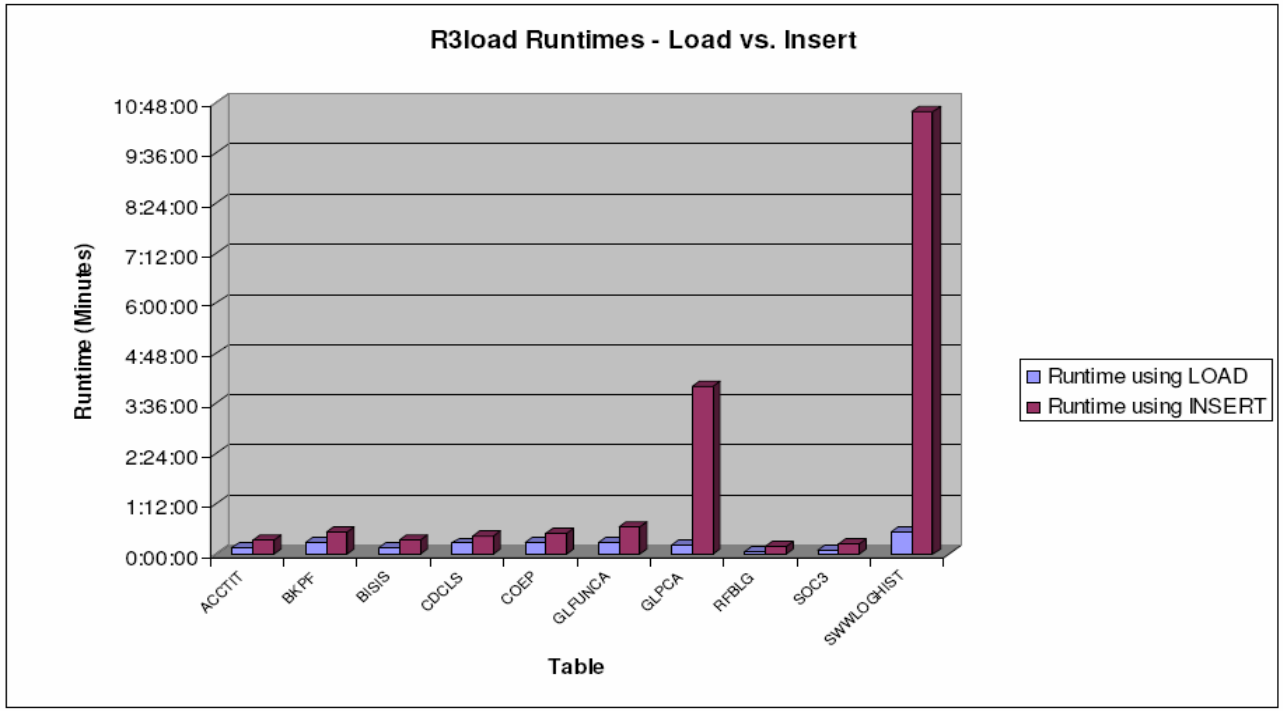

Testcase

- **Sequential import of tables**
- **Non-unicode**

Test results:

- Much better import performance with DB2 LOAD (optimized, parallel processing, reduced logging with DB2 LOAD)
- Some tables show less improvement with DB2 LOAD (CDCLS only factor 2 faster)

R3load invocation for optimized IMPORT

- DB2 LOAD can be invoked using the following R3load parameters

SAP Kernel 4.6D	`R3load -i ... -fast LOAD ...`
SAP Kernel 6.x and higher	`R3load -i ... -loadprocedure fast LOAD ...`

- INSERT with some optimizations will be used with the following R3load parameters:

```
R3load -loadprocedure fast
```

DB2 LOAD - Restrictions

INSERT is used by default, if:

- table contains LOB data
- table smaller than 200 KB
- table is splitted

DB2 Load will lock the table exclusively

- A splitted table (WHERE splitting) should therefore NOT be loaded with multiple R3load processes in parallel, when using DB2 LOAD
- The orderby-File (MigMon) can be used to define a package group with jobNum=1 for all packages of a splitted table.

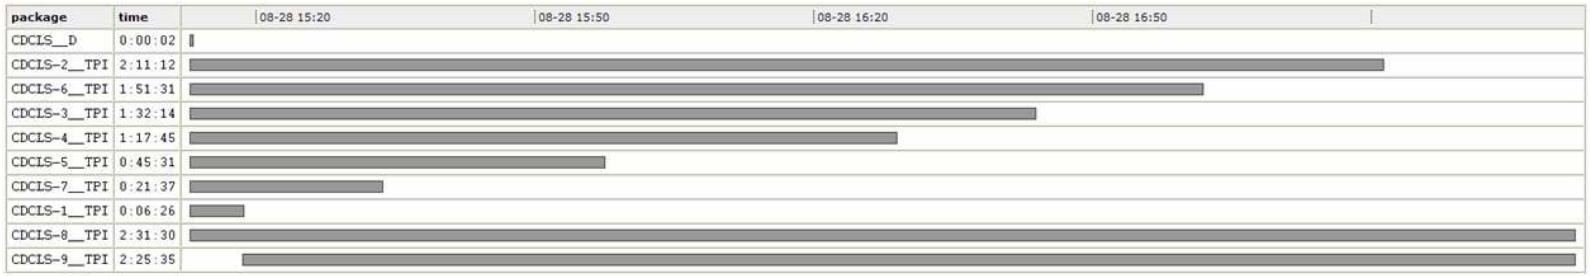

Import of splitted table with multiple R3load using DB2 LOAD

Important DB2 LOAD environment variables

DB6LOAD_CPU_PARALLELISM=<*n*> Defines number of processes or threads used by the LOAD utility to parse, convert, and format data records.

- The DB2 default is a good choice
- Monitor CPU parallelism (db2diag.log)
- If CPU is a bottleneck you may want to set this parameter to 2 or reduce the total number of R3load processes.

DB6LOAD_INDEXING_MODE =<*n*> Controls the index build: REBUILD (1), INCREMENTAL (2), DEFERRED (3), AUTOSELECT (0)

- The default (0) is recommended
- Exception: For splitted tables you should use INCREMENTAL mode (otherwise index would be rebuild several times)

DB6LOAD_FORCE_LOAD Forces DB2 to use LOAD even for small or **split tables** if it is set to any value (for example, 1). With R3load 7.xx, you can specify the force option in the fastload arguments on the command line.

DB2 LOAD vs. INSERT for Splitted Tables

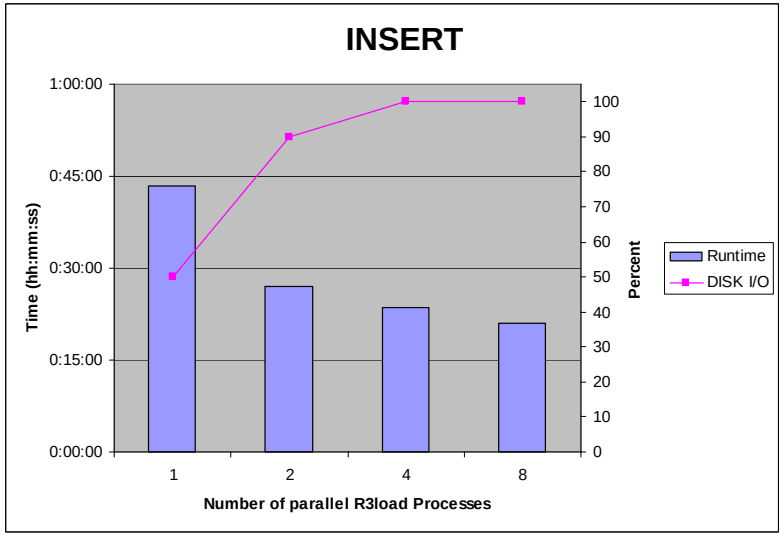

Parallel Import of Packages

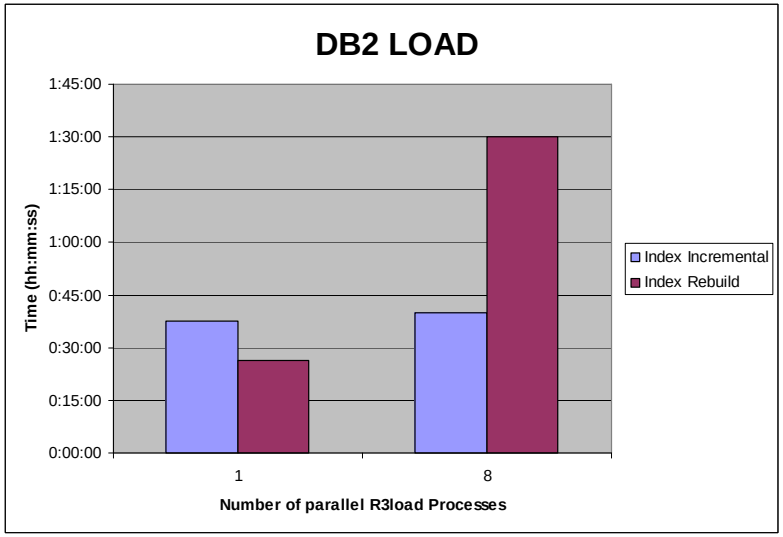

Sequential Import of Packages

Consideration **for splitted tables**:

- A single R3load process using DB2 LOAD may be faster than multiple R3load processes using INSERT.
- Multiple R3load processes using INSERT may require more hardware resources compared to a single R3load process using DB2 LOAD.
- For **UNICODE conversions**, the import time of a splitted table may be improved with multiple parallel R3load processes using INSERT. In this case the code page conversion is performed in parallel (R3load)

R3load invocation for splitted tables using DB2 LOAD

If you want to use DB2 LOAD for splitted tables (instead of INSERT):

- Use R3load option `fast LOAD_FORCED` (R3load 7.0) or environment variable `DB6LOAD_FORCE_LOAD`.
- Use a seperate MigMon instance to load splitted tables to make sure that small tables are NOT using DB2 LOAD.
- Make sure that packages for each splitted table are loaded sequentially by defining a package group for each table in the orderBy file.

```
[CDCLS]
jobNum=1
CDCLS-1
CDCLS-2
CDCLS-3
...
```

orderBy file

- Enable incremental index build for splitted tables.

```
#!/bin/sh
#
# Import Monitor startup on UNIX
#
export DB6_LOAD_INDEXING_MODE 2
...
```

import_monitor.sh

Creating Indexes before or after Import

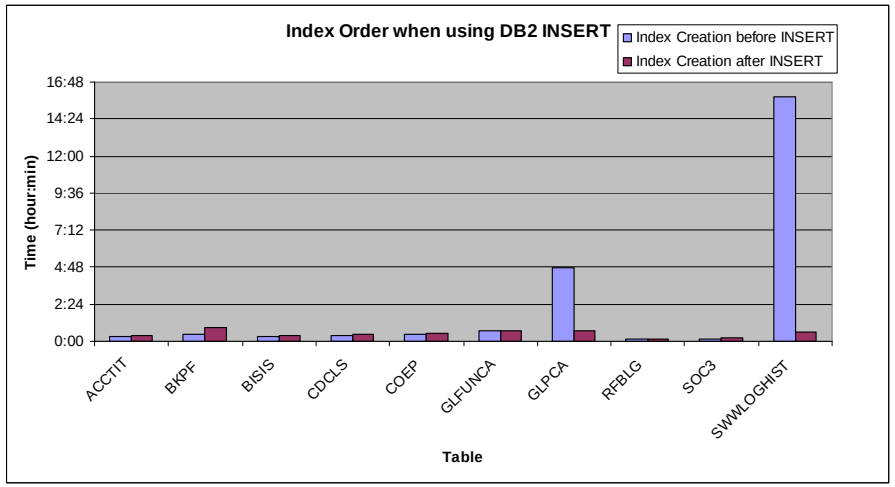

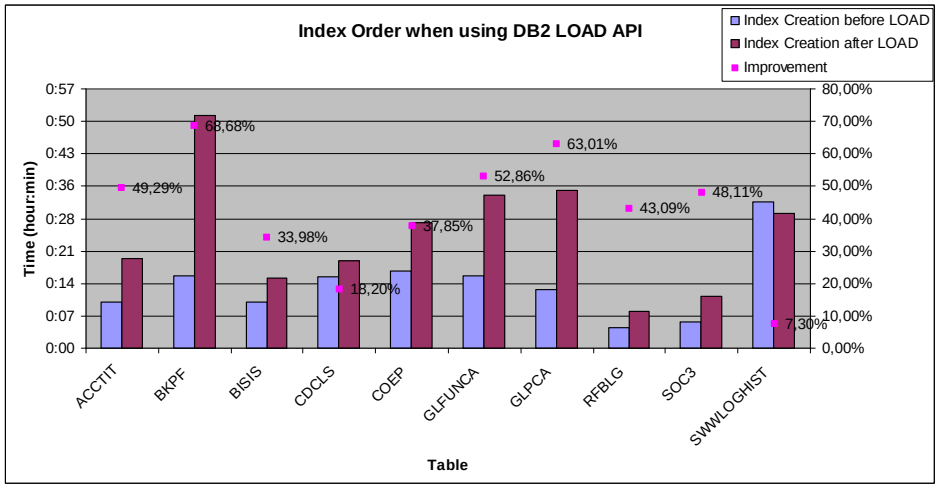

Test results with **INSERT**:

- In most cases index creation **before** import was faster.
- Only in few cases index creation after load was faster. In this case the difference was significant!

Test results with **DB2 LOAD**:

- In most cases index creation **before** import is faster.
- Only for table SWWLOGHIST index creation after load was slightly faster.

Indexes created before import are created during the *index build phase* of DB2 LOAD (optimized index build)!

Order of Index Creation - Conclusions

- Index creation before load is recommended (default since Kernel 7.00)
- In some cases creation of indexes after load may be faster. The difference can be significant. You may want to test both options for large critical tables!
- Ensure sufficient disk space for temporary table spaces. Tables and indexes may have grown between test and production migration.

- The index creation typically generates the highest I/O load during the migration process.
- To avoid I/O bottlenecks try to distribute the creation of indexes uniformly over the whole import time frame.

```
prikey:  BEFORE_LOAD ORDER_BY_PKEY
seckey: BEFORE_LOAD
cretab: CREATE TABLE &tab_name&
        ( /{ &fld_name& &fld_desc& /-, /} )
        IN &location&
        INDEX IN &locationI&
        LONG IN &locationL&
drptab: DROP TABLE &tab_name&
```

DDL<dbs>.TPL file

DB2 Optimizations for Import - Utility Heap

- UTIL_HEAP_SZ specifies the maximum amount of memory that can be used simultaneously by BACKUP, RESTORE and LOAD.
- If there is not enough Utility Heap available, the number of parallel processes used for a LOAD operation will be reduced.
- **Configure the Utility Heap at least to the number of parallel R3load processes multiplied by 30,000 pages**
- 200,000 pages is a good starting point if enough memory is available
- Overallocation of Utility Heap typically does not increase performance.

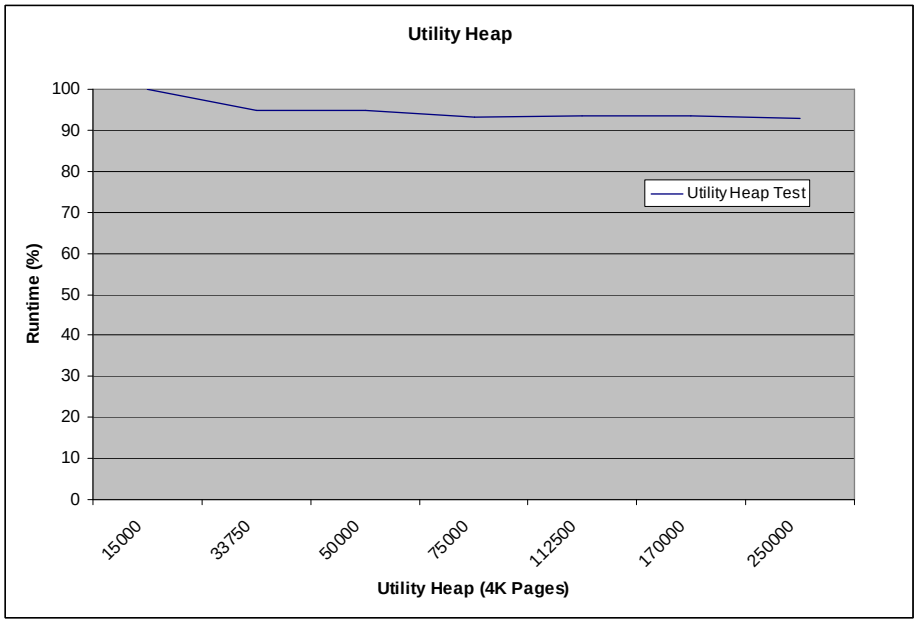

DB2 Optimization for Import - Sortheap

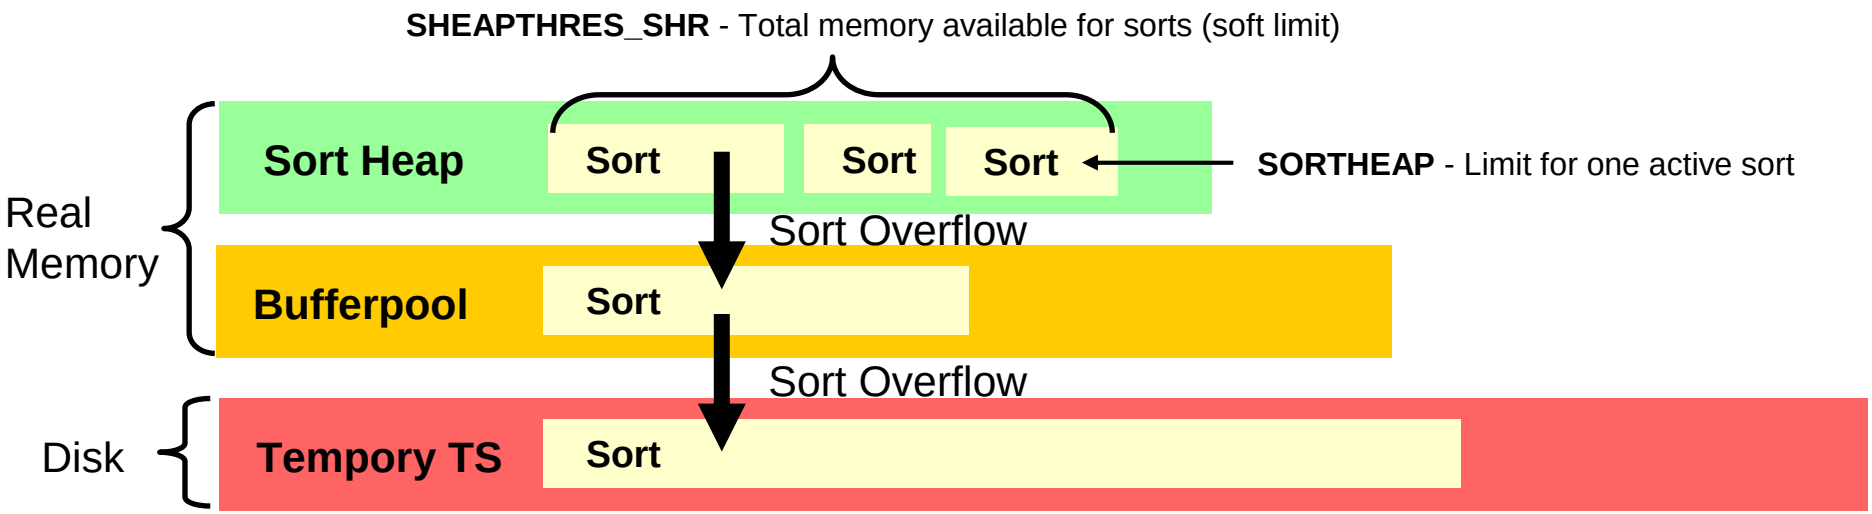

- **SORTHEAP** Define this parameter large enough to avoid sort overflows during index creation.

  A good starting point is 50,000 pages
- **SHEAPTHRES_SHR**

  A good starting point is 2 * SORTHEAP * (number_of_R3load_processes)
- SHEAPTHRES should be set to 0 to make sure that sorts are performed in shared memory.
- **To optimize sort parameters, monitor the database for sort overflows and adjust parameters accordingly.**

Number of R3load Processes

- As a starting point, use the <u>same number of R3load processes as CPU cores</u> are available on the server
- Monitor the system and make sure that all resources are optimally deployed during the entire migration

If the index build of many R3load processes overlap this can overload the machine (I/O) and degrade performance. In this case you may divide the import into two phases: (1) Data load, (2) Index creation

Migmon must be started two times with different configurations. For the index creation you can <u>lower the number of R3load processes, increase SORTHEAP, and enable INTRA_PARALLEL</u> (with DB2 9.5 INTRA_PARALLEL will be used by default to build indexes)

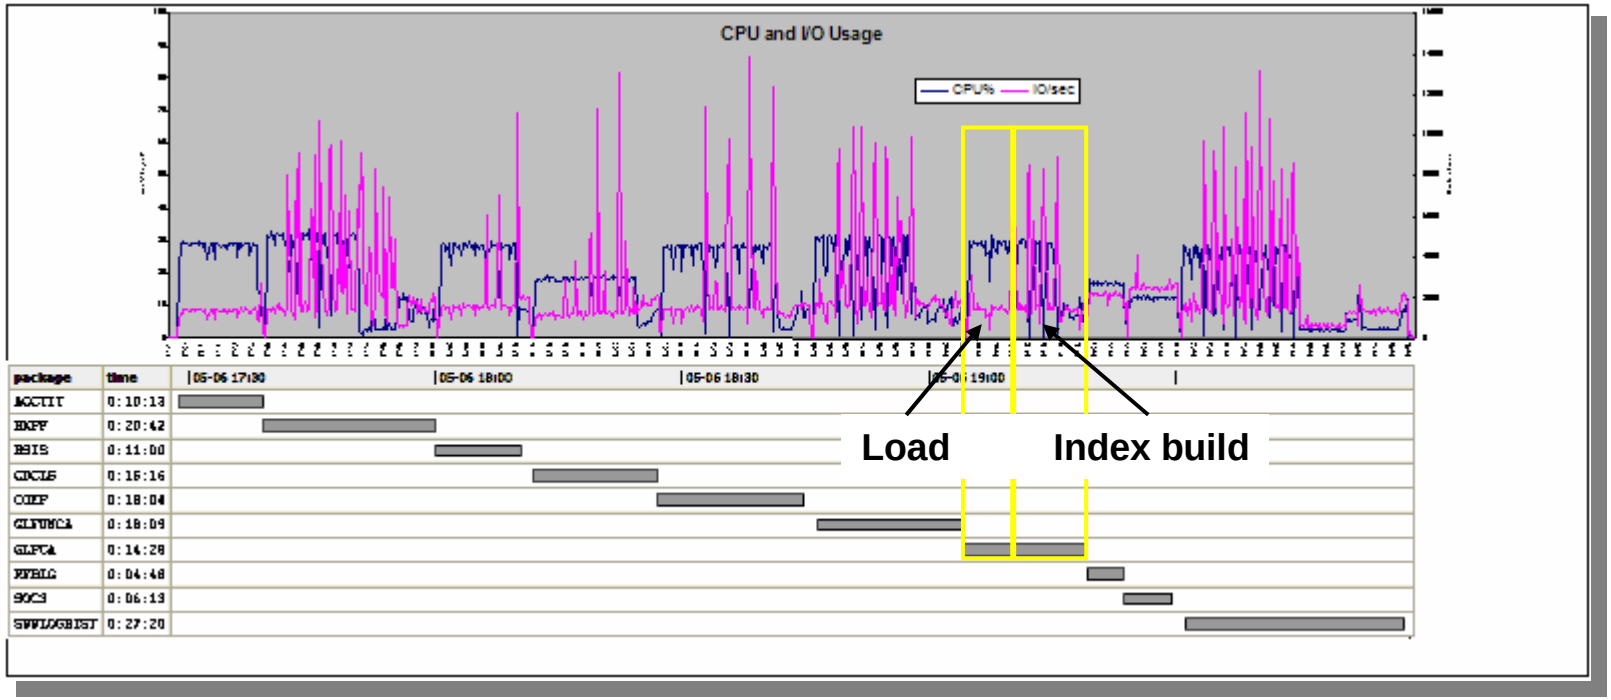

DB2 Compression

Name	Department	Salary	City	State
Miller	3875	10000	Dallas	TX
Smith	3875	15000	Dallas	TX
Johnson	4962	20000	San Francisco	CA

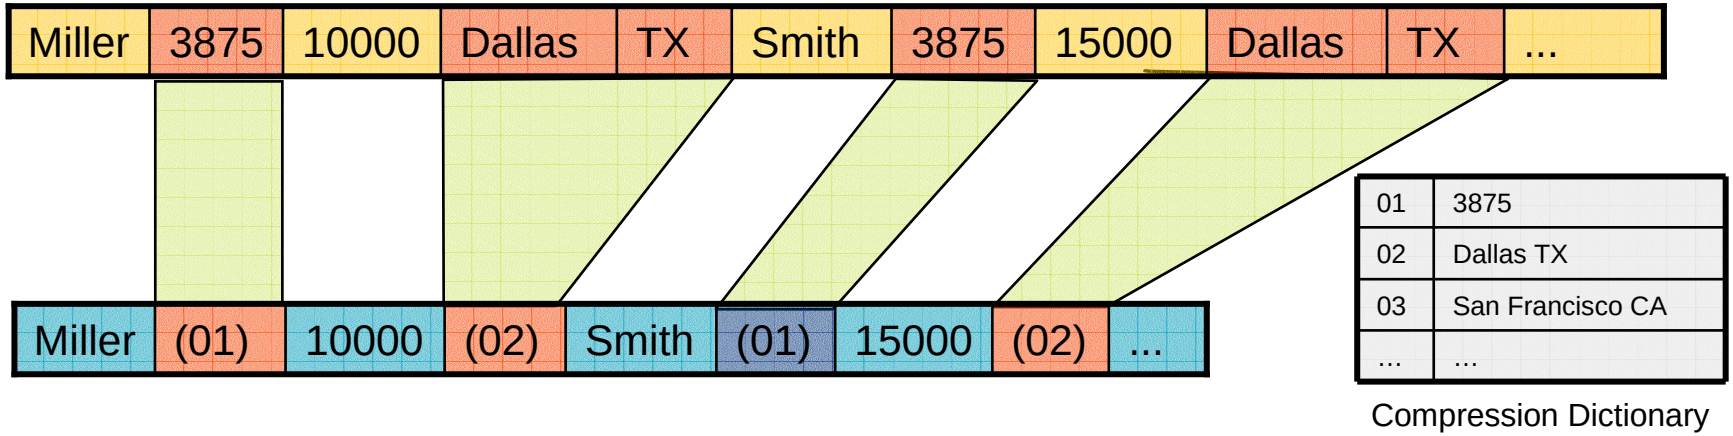

Compression Dictionary

Using compression:

- Enable compression flag: `CREATE TABLE <TAB> COMPRESS YES`
- Compress existing records (and optionally build new dictionary):

  `REORG TABLE <TAB> {KEEPDICTIONARY | RESETDICTIONARY}`
- If dictionary exists and compression is enabled new records will be automatically compressed.

Automatic Dictionary Creation (ADC) with DB2 9.5 if:

- Compression flag is enabled
- No compression dictionary exists
- Table contains sufficient amount of data (approx. 2 MB)

DB2 Compression based on R3load 7.00

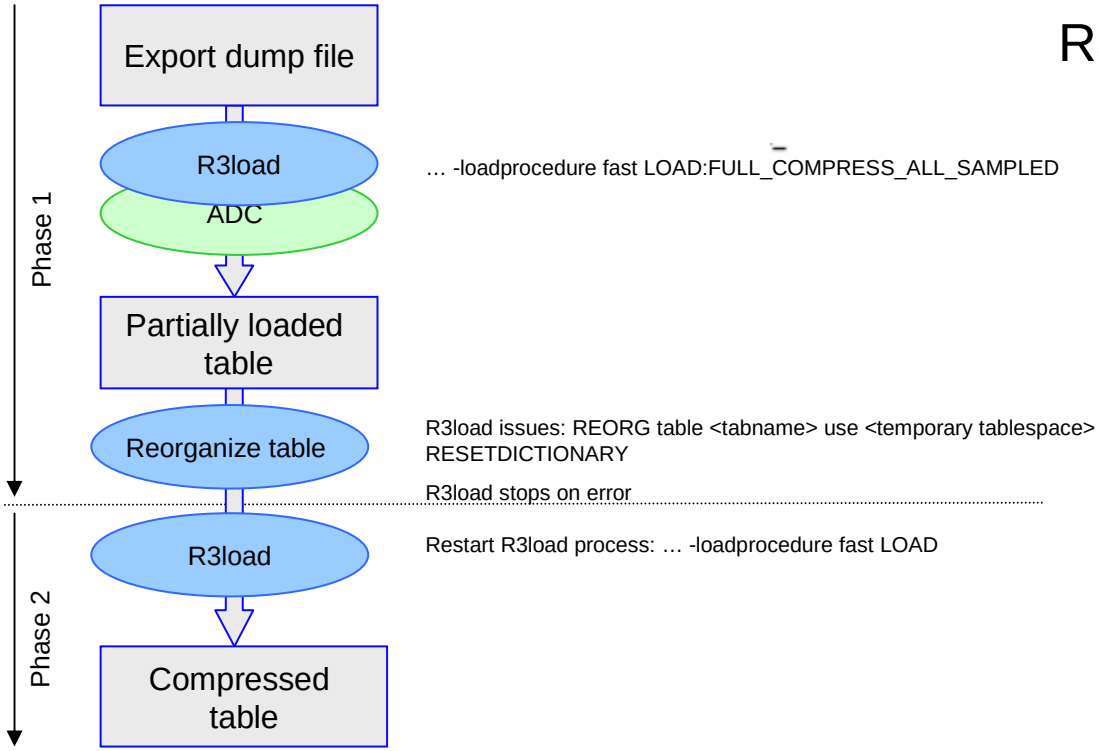

R3load must be started 2 times:

- **phase 1**
  - With option SAMPLED R3load imports a sample of the entire data.
  - Afterwards, it automatically performs a table reorganization with option "RESETDICTIONARY"
  - R3load stops with an error
- **phase 2**
  - R3load must be restarted without option COMPRESS

Advantages:
✓ Optimal DB2 compression ratios!
✓ No unused free space in the table space!

DB2 Compression based on R3load 6.40

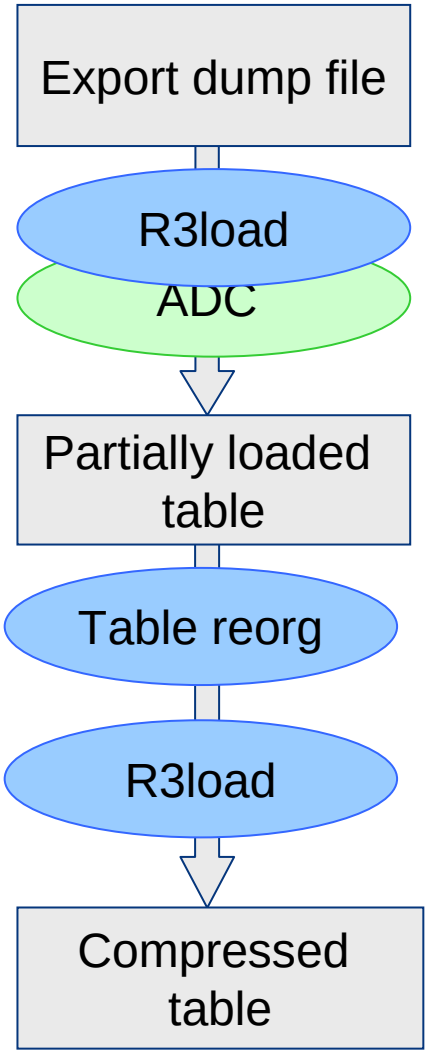

… -loadprocedure fast LOADCOMPRESS

Environment variable DB6LOAD_COMPRESSION_THRESHOLD defines the number of rows that are loaded before reorg is started (default 10.000). A sample of 1% of the table size is recommended.

R3load issues: REORG table <tabname> use <temporary tablespace> RESETDICTIONARY

R3load process continues to load data

Considerations:

- Compression ratios may not be optimal, because dictionary is based on subset of entire data.
- To further increase compression, additional reorg may be performed later on.

DB2 Statistics - Tuning Methods

- DB2 Runstats with option "SAMPLED"
  - Useful to improve DB2 Runstats performance for *large* tables
  - Generated statistics may not be optimal -> execution plans of SQL statements may not be optimal
- Parallel DB2 Runstats (multiple parallel scripts)
- Statistics creation with db2look:

```
select
'db2 "RUNSTATS ON TABLE <SAP-SCHEMA>.' || tabname
|| ' WITH DISTRIBUTION AND DETAILED INDEXES ALL "  '
from syscat.tables
where type='T'
and tabschema='<SAP-SCHEMA>'
and VOLATILE != 'C'   ;
```

Sample statement to generate DB2 runstats script

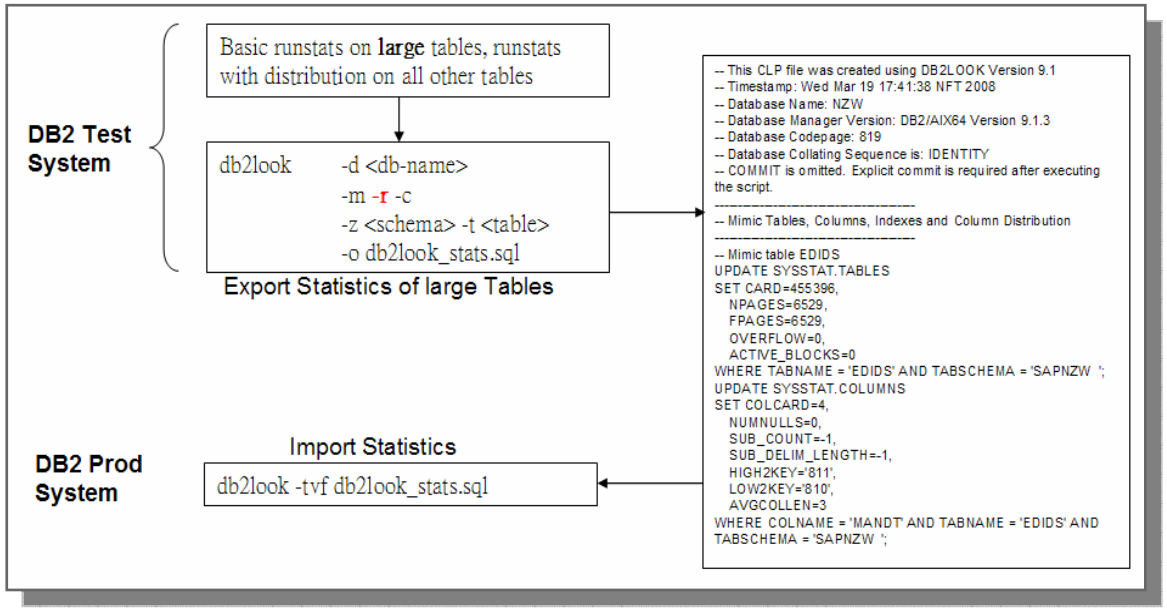

- db2look can only apply **basic** statistics if runstats was not executed before!
- Automatic runstats will be disabled for tables with imported statistics!
  To enable Automatic Runstats after import of statistics:
  db2 "RUNSTATS ON TABLE <tabname> SET PROFILE NONE"

DB2 Statistics - Summary

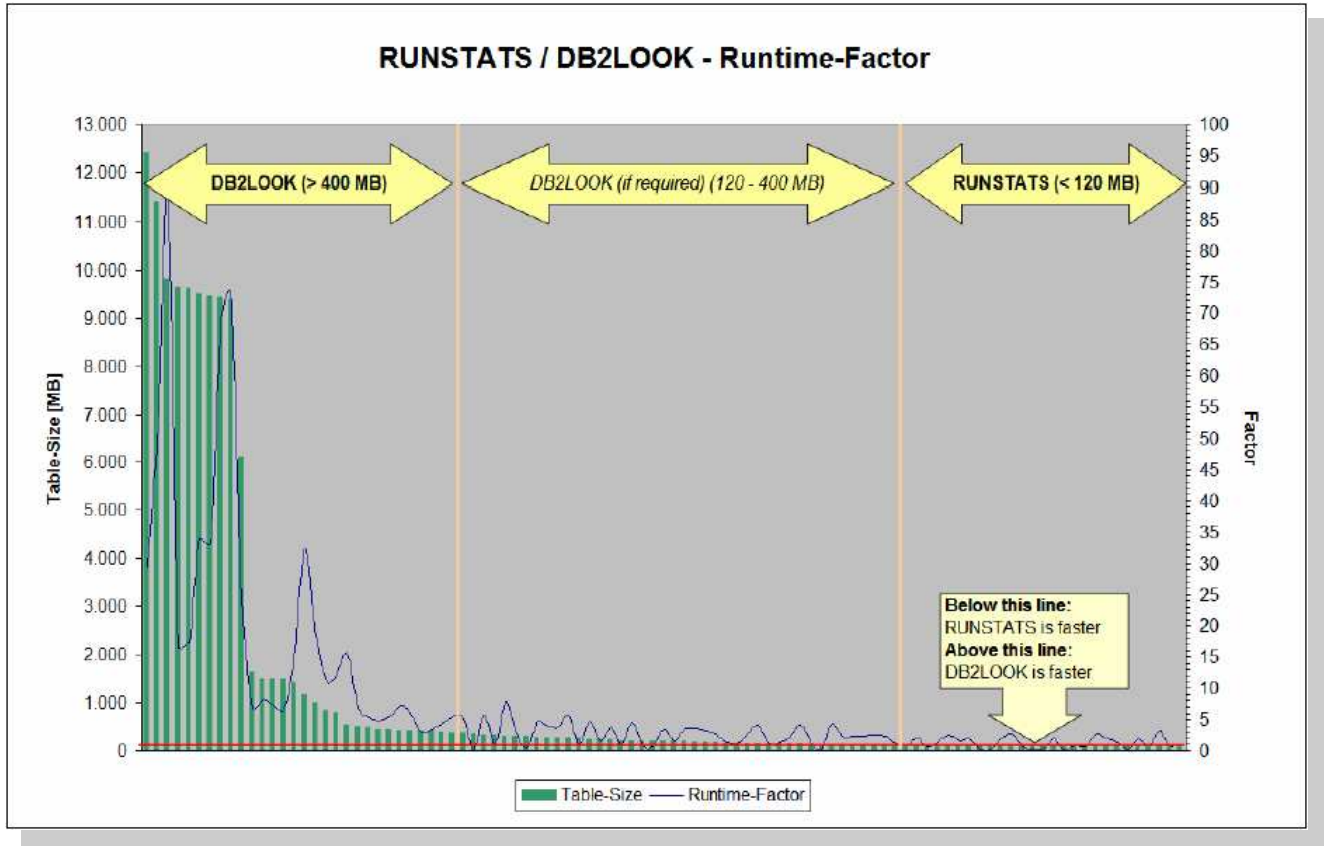

In our tests statistics creation with db2look was much faster than runstats for tables larger than **400 MB**

Statistics creation with db2look: Only runtime of generated script on target system was measured.

Summary

- Export Optimization
- Import Optimization

Migration Scenarios & Tuning Options - **Export**

Scenarios

Large Source Database
Impact:
- Long export times

Large Tables
Impact:
- Large tables dominate export runtime

UNICODE Conversion
Impact:
- Code page conversion is performed by R3load (higher CPU workload)

UNICODE Conversion with large table clusters
Impact:
- Additional processing and high workload in R3load for code page conversion
- Table clusters are always exported sorted!

Optimization Options

Package Splitting
- Reduce size of packages by splitting large tables into separate packages

OrderBy-File (Migmon)
- Export largest packages and tables first!

Table Splitting
- Export data of a single table with multiple R3load processes

Unsorted Export (TPL-File)
- Reduces export runtime
- Tables may need to be reorganized in target system (cluster ratio)

R3load on dedicated Server
- Adds additional CPU capacity
- Only useful if free I/O is available

Socket Option (Migmon)
- Reduces overall migration time
- Max saving is minimum of export or import time

Migration Scenarios & Tuning Options - **Import**

Scenarios

Large Source Database
Impact:
- Long import times
- Long database statistics runtime

Large Tables
Impact:
- Large tables dominate import runtime
- Long index build time, large sorts
- Long database statistics runtime

UNICODE Conversion
Impact of UTF16 to UTF8 conversion:
- INSERT: Code page conversion in R3load
- DB2 LOAD: Code page conversion in DB2

Compressed Target Database
Impact:
- Reduced I/O and therefore faster import

Optimization Options

DB2 LOAD Utility
- For large tables faster than INSERT
- Less resources required
- Table is locked exclusively

Parallel Import of Splitted Table
- Only possible with INSERT!
- UNICODE: Code page conversion in R3load (therefore parallel)
- With DB2 LOAD: Use serial import (define package groups) and incremental index build

Index creation before load
- Faster in most cases.
- With INSERT in some cases faster after load

Socket Option (Migmon)

Parallel Runstats or db2look

DB2 Compression with option SAMPLED (R3load)
- Optimal compression dictionaries
- No reorg required after import
- No High Water Mark issues

R3load on dedicated Server
- Adds additional CPU capacity
- Only useful if free I/O is available

Related Information

Training "Optimized SAP Migration to DB2"

- This 3 day workshop explains all optimization techniques for SAP Migrations with special focus on DB2 target systems. It shows detailed performance test results and the impact of the different tuning options
- Workshop participants will perform exercises on a VMware-based SAP/DB2 Demo Systems to gain hands-on experience with various techniques like table splitting etc.
- Contact: Liwen Yeow (yeow@ca.ibm.com)

service.sap.com/instguides

- NW 7.0: Installation & Upgrade Guides -> SAP NetWeaver 7.0 -> Installation -> System Copy
- NW 2004: Installation & Upgrade Guides -> SAP NetWeaver 2004 -> Installation -> SAP Web AS -> SAP Web AS 6.40 SR1 and Related Documentation
- Older Releases: Installation & Upgrade Guides -> SAP NetWeaver -> SAP NetWeaver Components (before SAP NetWeaver 2004)

https://www.sdn.sap.com/irj/scn/go/portal/prtroot/docs/library/uuid/c0546f90-998b-2a10-0fae-989576a8cb39

- "Heterogeneous SAP NetWeaver Business Intelligence (BI) System Copy to IBM DB2 for Linux,UNIX and Windows"

www.sdn.sap.com/irj/sdn/systemcopy

- SAP Guide "SYSTEM COPY & MIGRATION - OPTIMIZATION"

service.sap.com/osdbmigration

- Infos related to "SAP OS/DB Migration Check"

» Questions?

Thank you for your Attention!

Copyright 2007 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.

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

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

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

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

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

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

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

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

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

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

Alle Rechte vorbehalten.