Coca-Cola Bottling Co. Consolidated utilizes SAP technical upgrade project to migrate from Oracle to IBM DB2

About this paper

This technical brief describes the migration of an SAP R/3 Enterprise (version 4.7) non-Unicode system on Oracle Database 9i to a Unicode system with IBM DB2® 9 software. To achieve its business objectives, Coca-Cola Bottling Co. Consolidated (CCBCC) faced a technical upgrade of its SAP R/3 Enterprise system to SAP ERP 6.0. Completing this upgrade would require the company either to upgrade its existing Oracle database and purchase extra Oracle licenses or to move to a different database platform.

The CCBCC team decided it was time to derive more performance from the business-critical SAP applications while driving down hardware and software costs. Instead of upgrading Oracle, the team decided to introduce IBM DB2 software. As part of the SAP upgrade project, CCBCC's R/3 system would require a conversion to Unicode. By combining the Unicode conversion and the database migration, the client realized many technical and cost advantages—for example, exploiting time savings by sharing backup and test phases.

Client objectives

- Drive down software and hardware costs for the existing SAP software
- Complete the conversion and migration project on time and under budget
- Utilize the database migration project to avoid overhead for the Unicode conversion
- Gain better performance for the SAP applications
- Reduce storage needs
- Remove workload from the database administrators

"Running the SAP Unicode conversion and the database migration together produced a perfect result: the combination needed no additional downtime, and the project completed on time."

—Tom DeJuneas,
SAP systems manager, Coca-Cola Bottling Co. Consolidated

The solution

- IBM DB2 9.1 for Linux, UNIX and Windows technology, IBM DB2 Storage Optimization feature (Deep Compression)
- IBM Power Systems™ servers (IBM System p5® 560 model)

The benefits

- Combining the database migration with the SAP Unicode conversion saved time and money and caused essentially no overhead in effort for the database migration.
- Initial migration results show that even after the Unicode conversion, the DB2 software delivers a reduction in storage needs of approximately 40 percent as a result of the Deep Compression of the SAP R/3 Enterprise 4.7 system.
- The duration of manufacturing runs was reduced from 90 minutes to just 30—an improvement of more than 65 percent.
- The migration was completed under budget and ahead of schedule, with less than 26 hours of planned downtime—saving time and costs.
- The company has reduced overall licensing and maintenance costs by avoiding the purchase of additional Oracle licenses.
- DB2 software is easier to administrate and requires less attention from the SAP Basis database administrator—contributing to reduced costs.
- The company predicts savings in the next five years of about US$750,000.

Background, starting point and objectives

Coca-Cola Bottling Co. Consolidated makes, sells and delivers sparkling and still beverages, primarily products of The Coca-Cola Company. CCBCC is the nation's largest independent Coca-Cola bottler, operating in the United States, operating in 11 states, primarily in the southeast. Founded in 1902, CCBCC enjoys net sales of more than US$1.4 billion and is headquartered in Charlotte, North Carolina.

Leveraging synergies: SAP Unicode conversion and DB2 migration

Prior to the technical upgrade of the SAP landscape, CCBCC decided to perform both a Unicode conversion and a migration from the existing Oracle database platform to IBM DB2 software with Deep Compression. These changes would eliminate the need to buy new Oracle licenses and would thus reduce the total cost of ownership (TCO).

By switching on the DB2 Deep Compression feature during the migration, the company was able to reduce the size of the database by more than 40 percent—which is expected to result in faster backups and shorter run times for the upcoming SAP software upgrade.

In the meantime, before the SAP upgrade, CCBCC can benefit from the highly automated DB2 database administration, which offers reduced cost of operation. DB2 9 software includes features such as self-managing storage, self-tuning memory management (STMM), automatic reorganization, automatic runstats, real-time statistics and backup via the integrated IBM FlashCopy® feature.

All database administration and monitoring tasks can be completed from within the SAP Database Administrator (DBA) Cockpit for DB2 software—an easy-to-use management environment integrated into the SAP application environment.

Deploying Unicode as a future-proof solution

CCBCC decided to deploy Unicode because new SAP product releases (starting with SAP NetWeaver 7.0 onward) will be based on the Unicode standard. CCBCC wanted to be prepared for new SAP applications such as SAP NetWeaver Process Integration (SAP NetWeaver PI), which are already part of future implementation plans.

In technical terms, the requirements for a Unicode conversion are very similar to those of a database migration. In both scenarios, the client must perform an export and import of the database using the SAP program R3load.

The Unicode conversion itself is executed during the export phase of the migration. It is therefore very easy to direct the database toward a new target system without additional effort and downtime. Migrating to IBM DB2 software in conjunction with an SAP software upgrade and Unicode conversion leverages an opportunity to avoid duplicating project tasks such as backup and testing and keeps the cost of the migration as low as possible.

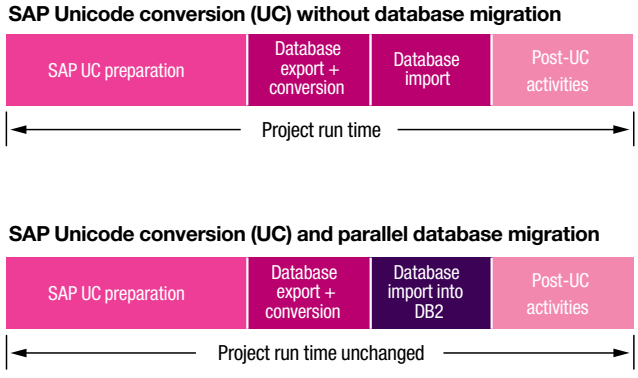

Figure 1: Combined database migration with SAP Unicode conversion

Migration process—heterogeneous system copy

CCBCC used a standard SAP methodology for the migration process, known as the heterogeneous system copy (or OS/DB migration) method. CCBCC was able to perform the migration and conversion during a scheduled maintenance window, so there was no need to make use of enhanced migration tools or services from SAP such as Zero Downtime.

The migration project for the entire SAP R/3 Enterprise landscape took eight weeks in total, including two test iterations for the 1 TB production database. The migration of the production SAP system itself was completed over one weekend, starting on the Saturday night and finishing in the early hours of Monday morning. The total downtime for the production migration was just 26 hours.

To achieve this reduced downtime, a set of SAP-specific migration tools were used:

- Unsorted export for the transparent tables
- Package Splitter for the largest tables ("big tables" group)
- Table Splitter for three large cluster tables
- Multiple instances of Migration Monitor to allow distributed parallel import and export processes
- R3load with Deep Compression option to activate compression during the migration phase

The next part of this document depicts the way CCBCC utilized these tools, explains the reasons for the choices and highlights the benefits.

Architectural overview—migration project at CCBCC

For the migration, CCBCC used four logical partitions (LPARs) on an IBM Power Systems server (System p5 560). Three LPARs were used to handle database export processes from the source system, and one LPAR was running the target system for the import processes. The export partitions consisted of a central instance, database (CI/DB) partition, which had 16 CPUs of 1.5 GHz and 64 GB of memory (CI/DB) and two other partitions that had four CPUs of 1.5 GHz and 12 GB of memory each. The import partition (or new CI/DB partition) had 16 CPUs of 1.5 GHz and 64 GB of memory.

During the testing phase, this system setup emerged as the optimal migration environment to handle the migration workload.

In order to meet the downtime objectives, the workload of the export packages were distributed between the CI/DB server and the other two servers (Hosts A and B) running in the first three LPARs. The CI/DB server handled the three largest cluster tables via Table Splitter. Host A handled the smaller tables. Host B was used to handle the export of the "big tables" group (which contained more than 10 million, more than 2 million and more than 200,000 records); these were divided into smaller packages using Package Splitter. All three hosts used local storage to dump the export data to disk. Each export process was controlled by a Migration Monitor (MigMon) instance with its own configurations.

On the import side there was only one server—Host C (new CI/DB server). The export disks of CI/DB, Host A and Host B were mounted via Network File System (NFS) (for reading) on Host C. The import was controlled by multiple MigMon instances.

From the "big tables" group on Host B, a subset was exported using the sorted unload option, which required additional CPU power and was one of the reasons for assigning an additional server for the export phase. During the import, the tables from the "big tables" group were compressed during the load process.

Database export—migration tools used unsorted versus sorted export

CCBCC used both sorted and unsorted exports to unload the data from the Oracle database. In general, the unsorted export is faster than the sorted. But as CCBCC was also running a Unicode conversion, the migration team was forced to export the SAP cluster tables (for example CDCLS, RFGLG, EDI40) and SAP repository data classes via a sorted export. Sorting the data required additional CPU power, which was one of the reasons CCBCC handled the export phase with three servers.

- Sorted export—pool tables, cluster tables, reports, Dynpros and nametabs.
- Unsorted export—most of the transparent tables

With a sorted export, the pages of a table are read in the sequence of the primary key. If the cluster ratio is not optimal, data pages will not be read continuously. In addition, database sort operations may occur, which will also extend the export run time. By using the unsorted option, data is read sequentially and written directly to a file instead of using an index that attempts to sort the data before writing to the file.

Unicode considerations for cluster tables

As a result of the Unicode conversion, the contents and the length of the records may change. Even the number of the physical records belonging to a logical record may change. Because the physical records are built together to form a logical record, the data must be read in a sorted manner to find all physical records that belong to a logical record. For these reasons, an unsorted unload is not possible.

Database limitations

DB2 software supports unsorted exports, but some other databases only allow sorted exports. This represents a major roadblock in migrating away from these databases and can also be a limitation in daily operations—for example, it is more difficult to set up test and quality assurance (QA) systems using sorted exports. Especially for very large databases, being forced to run a sorted export will heavily extend the downtime window and make it almost impossible to change the database or even complete a Unicode conversion in a reasonable time.

Package and table splitting

The database size of nearly 1 TB and the very large tables had been the determining factors for the downtime. CCBCC decided to parallelize the database export to improve the speed of the whole migration process by using Package Splitter and Table Splitter tools.

Package Splitter splits tables of the source database into packages and exports them. In each case a dedicated R3load process handles each package. These processes can run in parallel and consequently make better usage of the CPU power. Table Splitter R3ta generates multiple WHERE conditions for a table, which are used to export the table data with multiple R3load processes running in parallel. Each R3load process requires a WHERE condition so that it can select a subset of the data in the table.

- 262 large tables (“big tables” group) were put in their own package using Package Splitter to increase parallelism and ensure better granularity of the packages, resulting in better resource usage during the migration.
- 12 very large tables were divided into multiple packages using Table Splitter, enabling multiple R3load processes for parallel export and import of the table.
- The remainder of the tables were combined in joint packages using Package Splitter. By splitting the content to multiple R3load processes (20 parallel processes) it was possible to export and import the data in parallel, saving considerable time.

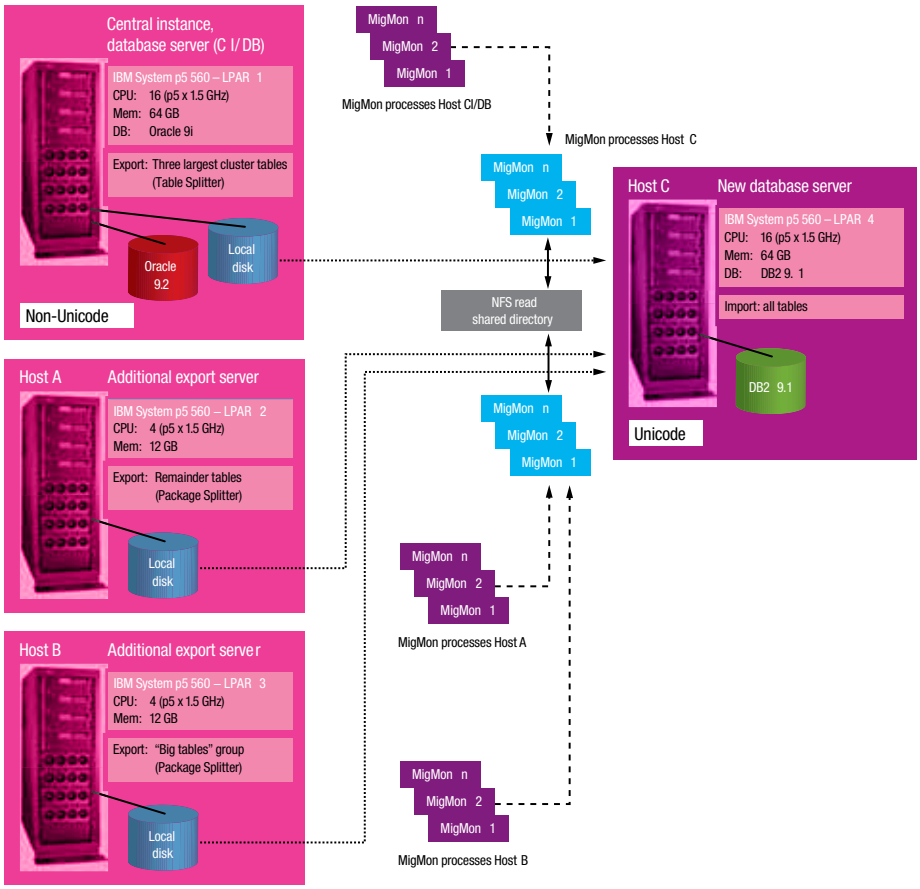

Figure 2: Landscape overview

Source:		Target:	
SAP release:	SAP R/3 Enterprise (version 4.7)	SAP release:	SAP R/3 Enterprise (version 4.7)
OS:	IBM AIX 5.3	OS:	IBM AIX 5.4
Database:	Oracle 9.2	Database:	IBM DB2 9.1
Database size:	950 GB	Database size:	575 GB (with compression)
Data format:	Single code page system without Unicode	Data format:	Unicode

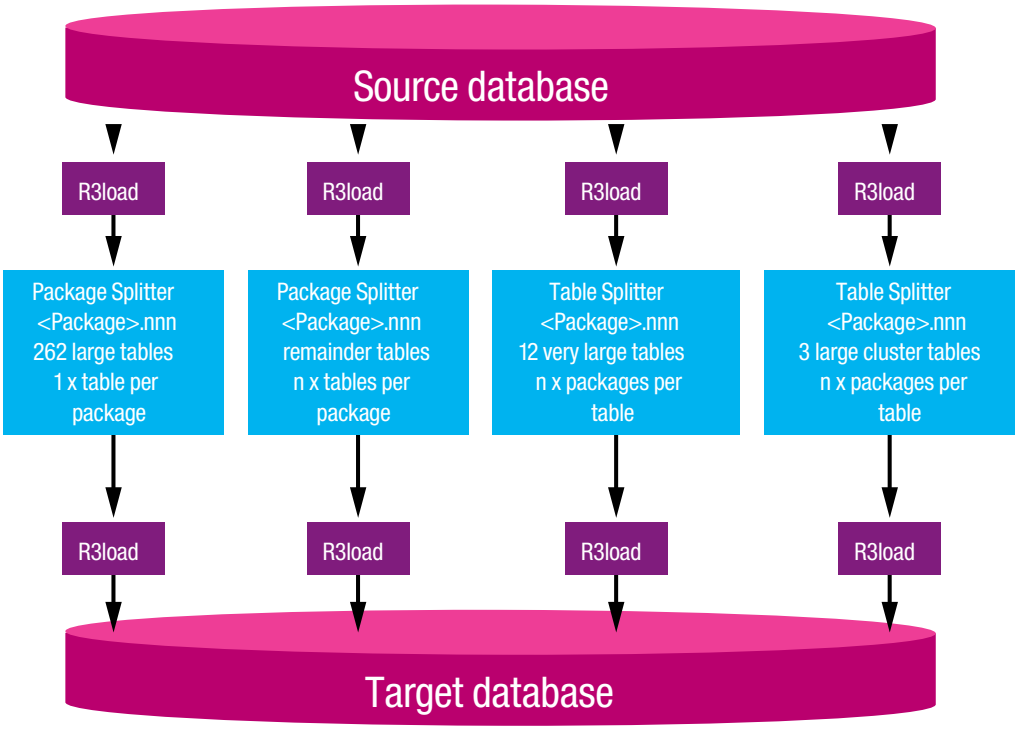

Figure 3: Package and table splitting

Migration Monitor

In a Unicode conversion, the system copy causes very high CPU load during the export. Most of the CPU power is spent on data conversion, especially when processing cluster tables. To avoid CPU bottlenecks, CCBCC distributed the exports and imports across four LPARs to parallelize these processes more effectively. This allowed CCBCC to take advantage of additional processor resources for the database export and import. The MigMon helped perform and control the unload and load process during the system copy procedure and enabled 20 export and import processes to be run in parallel.

Database import—DB2 Deep Compression enabled DB2 9 Storage Optimization feature

The DB2 9 Storage Optimization feature—also called Deep Compression—uses a dictionary-based approach to replace repeating patterns with short symbols. The dictionary stores the patterns that occur most frequently and indexes them with the corresponding symbols that are used to replace them. Due to the fact that all patterns within a table (not only within a single page) are replaced, impressive compression rates can be achieved (up to 90 percent for single tables).

"By choosing to implement DB2 compression right away, we have reduced the database size by around 40 percent. This gives us faster backup and reduced storage costs and makes the SAP technical upgrades easier and quicker."

—Andrew Juarez,
AP Basis lead, Coca-Cola Bottling Co. Consolidated

R3load with DB2 Deep Compression

CCBCC wanted to make use of the benefits that the DB2 Storage Optimization feature offers right away and decided to switch on Deep Compression during the migration process. Even with the knowledge that the compression rate with R3load version 6.40 might not be optimal, CCBCC decided to go ahead and was rewarded with a compression rate of 40 percent and an impressive performance improvement. This was achieved despite the fact that only the 169 of the larger tables had been compressed.

Enabling DB2 Deep Compression during database migration and Unicode conversion is a very smooth way to compress the data at the time it is loaded into the database. The R3load tool provides several ways of deploying DB2 Deep Compression when the data is loaded into the tables. Depending on the version of R3load (that is, version 6.40 or version 7.00 or higher), different options for compression are available, such as the new R3load 7.00 SAMPLED option.

This offers optimal data compression while avoiding time-consuming table reorganizations. In this paper we will focus on the compression feature of R3load version 6.40, as this was the tool used by CCBCC.

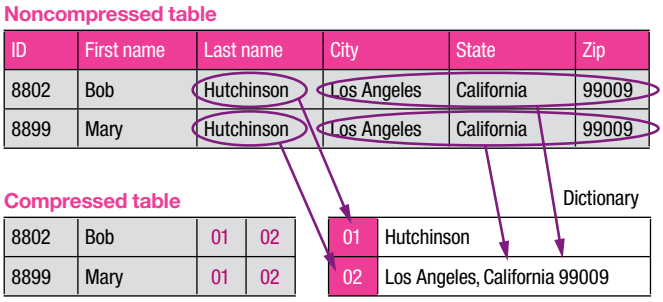

Figure 4: Deep Compression

R3load 6.40 with compress option

To generate the compression dictionary, R3load first loads a defined number of rows into the table without compressing them. R3load creates the compression dictionary based on these rows by running an offline reorganization.

CCBCC incremented the value of the environment variable DB6LOAD_COMPRESSION_THRESHOLD to define the number of rows that would be initially loaded and used to create the dictionary. The default value for this threshold is 10,000 records, which was too low to provide optimal compression sampling for the larger tables.

By sampling between 10 and 80 percent of the records (depending on the number of rows in the tables), CCBCC was able to set optimal threshold values and achieve very good compression results. The two largest tables (COEP, BSIS) contained more than 130 million records, followed by several tables with between 10 and 70 million records.

CCBCC grouped the compressible transparent tables using the following row count thresholds:

- Group of 20 tables of more than 3 million records
  – threshold = 3 million
- Group of 47 tables of more than 200,000 records
  – threshold = 200,000
- Group of 102 tables of more than 60,000 records
  – threshold = 60,000

Note that not all tables matching the thresholds were flagged for compression and added to those groups. Only the ones that showed good compression results in the test phase were selected.

After the initial import and the creation of the dictionary, R3load imports the remaining rows into the table, and the DB2 software compresses the data based on the dictionary.

Tables that are intended for compression during the load phase must have the compression attribute switched on. Since CCBCC had some tables that should be compressed and others that should not, different template files for the Migration Monitor were used.

CCBCC ran the import with several instances of the Migration Monitor (in different directories—see figure 2) and used different values for DB6LOAD_COMPRESSION_THRESHOLD for each instance.

Summary

Combining the Unicode upgrade with a database migration paid off for CCBCC, enabling the company to leverage synergies throughout the whole migration process and eliminate the duplication of processes such as backup and testing. The whole SAP ERP migration project took about eight weeks from start to finish, including the Unicode conversion.

Another essential aspect was the easy transfer of database management skills from Oracle to DB2 software and the user-friendliness of DB2 software. CCBCC had strong in-house Oracle skills, and yet in a matter of weeks the database administrators became fully competent on DB2 software, a tribute to the ease of transition to DB2 software for experienced DBAs, regardless of their technical legacy.

CCBCC was able to benefit right away from the value DB2 software offers:

- Lower TCO
- 40 percent reduction in database size
- Better performance—manufacturing runs are more than 65 percent faster
- Better integration of the database in SAP tools (SAP DBA Cockpit for DB2 software)
- Reduced DBA workload to manage and administrate DB2 software

With DB2 software in place, CCBCC is well prepared for the upcoming upgrade to SAP ERP 6.0, which can now be performed much more smoothly and rapidly. The reduction in database size by 40 percent will result in faster backup and shorter run times for the SAP software upgrade.

For more information about how you can lower your total cost of ownership while increasing performance, visit:

ibm.com/breakfreetoibm

IBM Corporation
Software Group
Route 100
Somers, NY 10589
U.S.A.

Produced in the United States of America
July 2011

Please Recycle