Success Story – Unicode Conversion with Incremental Migration (IMIG) at BASF

Company Name

BASF Corporation

Company Profile

BASF Corporation, headquartered in New Jersey, is the North American affiliate of BASF AG, Ludwigshafen, Germany. BASF employs about 10,000 people in North America and had sales of approximately $11.3 billion in 2005. For more information about BASF's North American operations visit www.basf.com/usa. As the world's leading chemical company, BASF's portfolio ranges from chemicals, plastics, performance products, agricultural products and fine chemicals to crude oil and natural gas. BASF's intelligent system solutions and high-value products help its customers to be more successful. BASF develops new technologies and uses them to open up additional market opportunities. It combines economic success with environmental protection and social responsibility, thus contributing to a better future. In 2005, BASF had approximately 81,000 employees and posted sales of more than $50.4 billion. Further information on BASF is available on the Internet at www.basf.com.

The system being converted is hosted in NAFTA and supports all companies in the Americas as well as many companies in Europe. This upgrade was to make the former Americas system become the Global ERP system of BASF Group. The upgrade for this system was coordinated with the central staff located in Rockaway, NJ.

Industry

Chemical

Key Challenges

- BASF was the first customer to use IMIG to convert their MDMP system to Unicode (treated as single code page conversion).
- The productive database of 4.5 TB was the largest database to be converted so far.
- Three systems – development, QA and production – with a database size of 2.6, 2.8 and 4.5 TB had to be copied. For each system a complete downtime (including all post-system-copy activities) of less than 36 hours was defined.
- The customer's production database contained five tables between 50 and 125 GB.

Product

- SAP R/3 Enterprise 4.70 with SAP R/3 Enterprise Extension Set 1.10 SR1 with add-on discrete industries. The system is also integrated with other SAP systems (mySAP CRM, SRM, SCM, BI and SAP Solution Manager).

Source System

- Database Server: AIX 5.2, 64 bit
- IBM p690 LPAR
- EMC Symmtrix 8830, 8 TB
- Oracle 9.2.0.7
- 40 GB memory

Target System

- Database Server: AIX 5.3, 64 bit
- IBM p5-590 LPAR
- Oracle 9.2.0.7
- 70 GB memory

Database Key Figures

- Database size (production): 4.5 TB (2.7 TB table data, 1.8 TB indexes)
- Largest tables:
  - DRAO (124 GB)
  - CKIS (103 GB)
  - COEP (67 GB)
  - BSIS (59 GB)
  - CDCLS (51 GB)
  - Growth of the database: 100 GB / month

Why did BASF choose IMIG?

Planned downtime for this system is 32 hours a year. BASF has much integration with scanners, plant systems that control manufacturing (DCS/PIM's) as well as weigh scales that control the flow of trucks at their plants. As such downtime is disruptive to their business processes and extremely costly to their businesses. Downtime was a critical factor in the Unicode conversion and an initial target of 24 hours downtime was established.

Procedure

The project was initiated one year before the final system copy of the productive system had to be terminated.

Three systems had to be converted to Unicode. With one additional test course a total of 4 IMIG runs had to be done:

1. As a test, an IMIG system copy of an (older) copy of the productive system.
2. Unicode conversion of the QA system.
3. Unicode conversion of the development system.
4. Unicode conversion of the productive system.

Test conversion:

The first test run took about 4 weeks. The largest 197 tables of the system had been chosen as IMIG tables, which covered ~88% of the database content. The largest 30 IMIG tables plus 8 IMIG tables which usually have long system copy runtimes too, had been split into smaller packages to enable parallel processing. In addition, 12 non-IMIG tables were split.

The main aim of this first test run was to validate the IMIG procedure and to verify that the IMIG does not introduce any errors when converting data to Unicode.

Results/Improvements:

- Unicode conversion with IMIG works fine.
- The number of IMIG tables could be reduced to 173.
- Some of the tables selected for splitting should be split into more and smaller packages and some additional table should be split to enable a larger overlap between export and import processes and thus reduce the total runtime. This resulted in 41 IMIG tables and 5 non-IMIG tables to be split.
- The packages of non-IMIG tables should be split into smaller pieces.
- Export and import order of packages should be changed to start with the long running packages and according to the additional defined packages (see above).
- A weak network cable led to serious i/o performance problems. After replacing the cable the issue was fixed.
- The cluster tables KOCLU and RFBLG contained invalid data which could not be imported. Tables were cleaned up in all systems.
- New features like unsorted unload, nologging and compute statistics during index creation and the Oracle fastloader should be used to further speed up the procedure.

Unicode conversion of the productive systems:

The results of each previous conversion were used to configure the Unicode conversion of the next system. Only few and small changes like modifying the number of packages or adapting the export and import order was necessary for the next IMIG run.

Using the Oracle fastloader reduced the import time by nearly 50%. The limiting factor is the time required for the index creation.

The key figures are summarized in the following table:

	BC2 – QA	BC3 – DEV	BC1 - PROD
DB size [GB]	2800	2600	4500
# CPUs exp/IMIG tables	10 @ 1.5GHZ	24 @ 600MHZ	20
# CPUs exp/non-IMIG tables	10 @ 1.5GHZ	24 @ 600MHZ	26
# processes exp/IMIG tables	15	12	12
# processes exp/non-IMIG tables	24	24	50
# CPUs import	10 @ 1.5GHZ	10 @ 1.5GHZ	12 @ 1.5 GHZ
# processes import	20	20	40
total project runtime [dd]	12	14	15
export IMIG tables [d:hh]	4:15	4:10	4:14
import IMIG tables [d:hh]	4:19	4:09	4:13
copy IMIG tables [d:hh]	4:20	4:10	4:19
export non-IMIG tables [hh:mm]	11:00	07:10	4:20
import non-IMIG tables [hh:mm]	11:30	07:00	6:15
copy non-IMIG tables [hh:mm]	12:00	07:50	7:00
total downtime* [hh:mm]	27:00	20:00	16:00

* total downtime included: Unicode specific post-unload activities, IMIG specific downtime preparations, unload/load, post-load activities, Unicode specific post-load activities, update database statistics, report generation, system checks and reconfiguration, backup.

Implementation Highlights

- Just one pure test run required! No long-winded testing/configuring cycles necessary to reach the given downtime frames. Each IMIG run contributed to the improvement of the next IMIG system copy.
- Means to increase the performance were uncomplicated (some more splitting, changing processing order, using fastloader) and straight forward. No sophisticated knowledge of OS/DB tuning was necessary.
- The productive system could be released ahead of schedule after a downtime of only 16 hours! The downtime included 7 hours for unload/load data, 4 hours for updating the database statistics, 4 hours for generating the report load, 2 hours for performing the post Unicode conversion steps and time for reconfiguring and testing the new Unicode system.
- Virtually no issues after the conversion.

Key Benefits

- Due to the very nature of the IMIG procedure, only a small part of the database had to be copied within the given downtime window. Thus the customer requirements were met very easily.
- Only 16 hours downtime (including all post-load activities and system checks) for copying a 4.5 TB database!