



CDB Software, Inc.

VSAM to DB2 Conversion

June 27, 2006

Introduction

Much has been written about why you might want to convert your VSAM applications to DB2. Application development, by and large, changed from flat-files or files with an Access Method (like VSAM) to Data Base Management Systems (DBMS) like DB2 about 30 years ago. DBMSs brought us lock managers that allow concurrent access to your data, even from batch and online applications concurrently. DBMSs enhance security with the ability to control access at the record and field level and audit such access. DBMSs enhance recovery with the ability to back out and redo units of work in sync with other databases, other DBMSs, and Message Queuing systems. DBMSs, specifically DB2, is scalable with tables and indexes that can each be multi-terabyte (1000's of GB) with each object concurrently available from multiple Mainframe computers. However, the big draw to Relational DBMSs like DB2 is Accessibility. DB2 data is available to everything from big Data Warehouses and Data Marts for Analytical processing, IMS and CICS for transaction processing, Client/Server, Network, Desktop and Mobil Computing applications for Decision Support and everything else. This makes for more productive users as your data becomes information available to reporting applications and spreadsheets on desktops. And developers are more productive when much of the reporting requirements are fulfilled by end-users with simple SQL based reporting tools.

But converting VSAM data to DB2 tables is not for everyone. If your company doesn't have DB2 skills and doesn't plan to develop them then this article is not for you. If your VSAM applications are running well, with no support staff issues, and serving your information requirements; leave well enough alone. If your VSAM application is purchased software supported by the vendor

and must be upgraded in a tight time schedule to maintain compliance with federal regulations, you might want to minimize custom changes. Finally, if you're already replicating your VSAM data to DB2 and don't mind the delayed availability or the resources expended to maintain a second copy of the data you may not be interested in this article.

Since you're still reading, you're interested in the benefits of converting VSAM to DB2 (or trying to figure out how to get that done more easily) so let's review your options for getting the conversion done.

Methods

The quick and dirty way is to make each VSAM dataset a DB2 table and convert your application's VSAM calls to DB2 calls. But, after recoding all your applications you will not have an optimally performing system. Also, you will not realize the accessibility gains you were looking for because the DB2 database is not normalized; i.e. redundancies eliminated, data organized efficiently, and data consistency improved. This brings us to the second option.

The thorough but longer method is to design a normalized DB2 database, then use utilities and home-grown programs to move your data to that structure. Then your applications must be recoded to access DB2 where they used to access VSAM. This will most likely leave you with an application system that will out-perform the original. It will certainly realize all the gains discussed earlier in the areas of Security, Recovery, and Accessibility. The down side to this approach is the size of the project which brings risk to its success. This method is an all-at-once cutover project that either freezes changes to the original system, or adds more complexity to the coding and testing. This brings



us to the new way and the title of this article.

The New Method

For the past five or six years there have been products that not only help you convert your VSAM datasets to normalized DB2 tables, they also allow your unmodified applications and VSAM utilities to access those DB2 tables. These new transparency tools change everything. They intercept VSAM calls, get the data from DB2 with SQL, and pass back to the VSAM application the data in the form it expects or the appropriate VSAM return code. Now you can convert VSAM datasets to their corresponding DB2 tables one at a time and this conversion is transparent to the application. Existing VSAM applications are converted to DB2 applications even if the source code no longer exist. You can convert one VSAM file to DB2 leaving the others in VSAM. Your applications will run untouched and access both the new DB2 table and the other existing VSAM files. The project involves many smaller successes over time rather than an all-at-once cutover. This allows you to start with small files to get your feet wet, and then tackle the more significant files. As soon as a dataset is converted to DB2 it is immediately available to existing database reporting tools and new DB2 program development. The original VSAM applications can be converted one at a time to native DB2 calls or left untouched to access DB2 through the DB2 transparency module. To make testing smooth, you can easily include and remove VSAM files from the transparency list. Also, you can define transparency for specific terminals or batch runs. Finally, the VSAM to DB2 tool can automatically re-migrate from DB2 to the original VSAM file if you desire to fall back. This method of VSAM to DB2 conversion using DB2 transparency greatly reduces the size and risk of the project and starts realizing benefits immediately.

Now for some detail about how these new VSAM to DB2 conversion tools work. VSAM comes in three forms, ESDS, RRDS, and KSDS. Each of these can have standard records, multiple record types, and repeating groups. The transparency tools should handle all these cases automatically, transparently, without the need to write any code or exits. They also include a migration tool that produces normalized DB2 data structures automatically. These VSAM to DB2 migration tools provide a TSO interface to map a VSAM file to DB2 tables. They automatically handle multiple record types with DB2 tables and views, create DB2 parent-child related tables for VSAM Occurs clauses, and calculate relative record numbers which translate to an indexed column in DB2.

Once the DB2 structure is defined the VSAM to DB2 migration tool needs to handle data conversions. This not only means date conversions and numeric conversions but also translating spaces or low-values to DB2 null values and translating special dates like all 0's or 9's to Null. When using the transparency tools, you use the TSO based tool to specify conversions for each field. This way the translation occurs efficiently along with the Unload and Load of the data to create proper DB2 data types. The Null values allow query tools to calculate the correct summaries and averages etc. The DB2 transparency modules translate these proper DB2 values back to what the original VSAM modules expect so that they can continue to execute properly. Failure to translate these data "special cases" in VSAM to Nulls in DB2 would result in tables with limited functionality going forward.

So your steps are: 1) Name the VSAM file 2) Define fields and data conversions (copy book record layouts are helpful here) 3) Click to



generate static SQL modules 4) With another click, DB2 table DDL is generated along with VSAM unload and DB2 Load statements. 5) Create those DB2 tables in your DB2 Databases and Tablespaces 6) Run the Unload and Load 7) From a TSO Panel, activate file transparency for this VSAM file.

Now all your application programs and VSAM utilities that access this VSAM file will get their data from DB2. Your applications will continue to get their other, not converted, VSAM data from VSAM. No changes to JCL, Batch code, or online code. The converted data is now available to report applications, spreadsheets, and new DB2 applications.

Summary

With your data in DB2 you get all the benefits of a Relational DBMS mentioned earlier without rewriting your original VSAM applications. However, those original VSAM applications could benefit from being rewritten to take full advantage of DB2. VSAM processes records while DB2 processes related sets of data. DB2 has clever ways to join together related tables and choose a path to the data independent of your program code. To take full advantage of this your programs would be rewritten with embedded SQL. This process of using DB2 transparency took a daunting task and turned it into easily manageable tasks with immediate successes

and benefits that end with a normalized DB2 database and the opportunity to write new or modify existing programs to take full advantage of it.

The savings in programming, consultants, and testing, along with the benefits of having your data in DB2, make this a cost effective solution with a very high probability of success.

Some of the products that make this conversion easier are: DB2 VSAM Transparency from IBM, VSAM2DB2 from XTISOFT (available through CDB Software), and VS/2 from Circle Computer Group. An evaluation of these products will show varying degrees of effectiveness.

Contributed by Doug Wilson

Doug Wilson has over 20 years of DBA experience with companies that included First Data, Prudential, Watkins Motor Lines and began with Lockheed Martin. He now works as a Technical Consultant for CDB Software, Inc.

About CDB Software

CDB Software, Inc. is a leader in data management solutions for DB2 z/OS. CDB focuses its business on DB2 for z/OS to provide unique and innovative solutions that enable companies to expand their DB2 system to meet business needs while controlling the overall cost of the mainframe. Founded in 1985, CDB is a privately held corporation based in Houston, Texas with offices worldwide.

For more information visit:

www.cdbsoftware.com



CDB Software, Inc.
11200 Richmond Ave.
Houston, TX 77082