



Jackie Jansen

Warehouse Management Tools

In the last few issues we have discussed what an informational database is and what tools you would use to access it. Today we are going to look at how you would build and manage this informational database.

If your plans are to build a single summary table or a single small datamart then quite honestly you can write your own program to do this. Coglin Mill suggests that their product Rodin is cost justified if you are planning on implementing three or more datamarts. At this stage the maintenance of your source and target data definitions becomes a significant concern.

Now this is where I start to get into trouble. There are a lot of tools available in the marketplace. I am going to discuss the three that I see most frequently implemented among iSeries customers. These are IBM's OLAP Builder which is based on ShowCase's Warehouse Builder, Colgin Mill's Rodin and IBM's Warehouse Manager.

Until recently there were some obvious differences between these products but some of the lines are starting to blur. From a high level point of view both OLAP Builder and Rodin are iSeries centric products with OLAP Builder being more of an entry level product and Rodin being a more robust product. Warehouse Manager is a product that will run across IBM's eServer platform and requires Microsoft NT as a base.

Both OLAP Builder and Warehouse Manager are SQL based products. Rodin is an ILE RPG generator.

In a larger environment this gives Rodin some real benefits. SQL restricts you to creating a maximum of one output table per pass of the source data. In addition you can require multiple passes of your source data to create a single output table. Rodin rarely needs more than one pass and can create multiple tables, including detail and summary tables, in the same pass. IBM internally uses Rodin in our teraplex centre when we publish our high end iSeries data warehouse load benchmarks. This product makes excellent use of an SMP or clustered iSeries environment and can have multiple tasks all running in parallel utilizing all your CPUs simultaneously.

All three products work very well to create a relational database. OLAP Builder and Warehouse Manager are nicely integrated with IBM's DB2 OLAP Server and ShowCase's Essbase/400. Rodin requires a user-written CLP to load a multidimensional database.


Until recently Warehouse Manager had a large advantage when it came to reading data from multiple different sources such as DB2/400, SQL Server, Oracle, MS Excel etc. OLAP Builder added this functionality earlier this year and Coglin Mill has announced that they will be able to read external data sources in their latest release that is due to ship at about the time you will be reading this article.

Both Warehouse Manager and Rodin support the creation and access of metadata. If you are running an ERP system you might want to check which products will read your data dictionary if that is a concern.

For example both OLAP Builder and Rodin can read a JDE data dictionary whereas Warehouse Manager understands the data elements for I2 and SAP.

In a smaller fairly homogeneous environment we often see OLAP Builder as the product of choice. It is iSeries or AS/400 centric, easily understands DB2/400 tables and files and works with both relational and multidimensional data.

Warehouse Manager, besides running on all IBM platforms, also has a set of "transforms" or pre-written routines that the user can invoke. One of these routines allows you to switch around rows and columns for data loading ie you can take 12 monthly sales buckets in 1 record and convert them to 12 individual records. There are transforms to build date tables containing almost any date attribute you can imagine and routines to perform mathematical and statistical functions.

I have just touched the surface here of the products available. All in all it should be a comforting feeling to know that you have choices that range from small, simple solutions to extremely robust industrial strength solutions, all running on the iSeries. 

Jackie Jansen is a Senior Consulting IT Specialist. She currently works in the IBM Americas Business Intelligence Solutions Centre. Jackie is a frequent speaker at iSeries Technical Conferences and User Group meetings. Contact her at jjansen@ca.ibm.com.