



Replicating CA's MDB Database

Published: 1st February 2006
Channel: Data Management

MDB stands for Management Database and it lives at the heart of Computer Associates' strategy for on-demand computing, which it refers to as Enterprise IT Management (EITM). This integrates more than 25 of the company's products, such as Unicenter, eTrust, and so on. In the middle of EITM is MDB, which is used to store and manage the information needed so that applications such as Asset Management or Service Desks can do their thing. Moreover, MDB provides data level integration between the various EITM-enabled products.



MDB can be based upon Ingres (by default) or SQL Server (and, in the future, potentially Oracle), which is all fine and dandy. However, it will often be the case that you will not want a single, monolithic MDB implementation. This might be for various reasons: you might want a back-up for failover and standby reasons, you might want to use a federated approach where you have an environment distributed across multiple time zones, or you might want to have one MDB for your Service Desk and another for Asset Management, for example (which is not at all unlikely given that a fully configured MDB can have upwards of 2,500 tables).

If it is the case that you have multiple or distributed MDBs then you will need some way to replicate data from one environment to another. You might also need this if you want to propagate information from an MDB to a data warehouse for analysis purposes. The question is: how do you achieve this?

One wouldn't think that this should be a problem, but it is: replication in Ingres is limited at best. And while it is better in Oracle and SQL Server the truth is that the replication in these databases was really designed (in the first place) to support remote salespeople and the like who logged into the central database once a day to synchronise their laptops. As a result, these products tend to implement replication on a table by table basis, which means that performance is poor; and they lack resilience. You also have to drag-and-drop tables on screen to implement replication, which is fine when you only have to replicate a few tables but won't work when you have hundreds, if not thousands as you often will with MDB.

The solution then, is a third party specialist replication. One such company is PSB, which is a Dutch company that started life as a specialist Ingres partner but which now supports both Oracle and SQL Server as well, with its HVR product. This stands for High Volume Replicator and it does what its name suggests: using streaming to provide high throughput and performance rather than using a table-by-table approach.

Further, the HVR product has a lot more scope for putting "intelligence"

into the replication. For example, it has the ability to recognise different versions of the databases it is addressing and it can replicate to different data models that may exist in different instances. This is particularly relevant in the distributed environments discussed above (as opposed to using the software for standby and high availability) because it will often be the case that, say, the Service Desk and Asset Manager are using different versions of the database. In addition, because it can recognise these differences, this means that HVR can be used in an elegant manner to support upgrades to the environment.

Historically, PSB simply sold and marketed replication as a technology. However, this limited its market to the very largest companies. The company is now more solutions focused so that it has specific offerings to support applications such as the Unicenter Service Desk, which may be implemented as point solutions. Whether you are just looking for an application such as this, or you want a broader solution, HVR is well worth a look.

Philip Howard
Director of Research, Bloor Research

 **Bloor
Research**
...optimise your IT investments
© Bloor Research 2006